

HUMAN
COGNITIVE
PROCESSING 24

New Directions in Cognitive Linguistics

Edited by
Vyvyan Evans
Stéphanie Pourcel

John Benjamins Publishing Company

New Directions in Cognitive Linguistics

Human Cognitive Processing (HCP)

Human Cognitive Processing is a bookseries presenting interdisciplinary research on the cognitive structure and processing of language and its anchoring in the human cognitive or mental systems.

Editors

Marcelo Dascal
Tel Aviv University

Raymond W. Gibbs, Jr.
University of California at Santa Cruz

Jan Nuyts
University of Antwerp

Editorial Advisory Board

Melissa E. Bowerman
Nijmegen

Wallace Chafe
Santa Barbara, CA

Philip R. Cohen
Portland, OR

Antonio Damasio
Iowa City, IA

Morton Ann Gernsbacher
Madison, WI

David McNeill
Chicago, IL

Eric Pederson
Eugene, OR

François Recanati
Paris

Sally Rice
Edmonton, Alberta

Benny Shanon
Jerusalem

Lokendra Shastri
Berkeley, CA

Paul Thagard
Waterloo, Ontario

Volume 24

New Directions in Cognitive Linguistics

Edited by Vyvyan Evans and Stéphanie Pourcel

New Directions in Cognitive Linguistics

Edited by

Vyvyan Evans
Stéphanie Pourcel
Bangor University, UK

John Benjamins Publishing Company
Amsterdam / Philadelphia



The paper used in this publication meets the minimum requirements of American National Standard for Information Sciences – Permanence of Paper for Printed Library Materials, ANSI Z39.48-1984.

Library of Congress Cataloging-in-Publication Data

New directions in cognitive linguistics / edited by Vyvyan Evans, Stéphanie Pourcel.

p. cm. (Human Cognitive Processing, ISSN 1387-6724 ; v. 24)

Includes bibliographical references and index.

1. Cognitive grammar. 2. Linguistics. I. Evans, Vyvyan. II. Pourcel, Stéphanie.

P165.N48 2009

415--dc22

2009010214

ISBN 978 90 272 2378 4 (HB; alk. paper)

ISBN 978 90 272 8944 5 (EB)

© 2009 – John Benjamins B.V.

No part of this book may be reproduced in any form, by print, photoprint, microfilm, or any other means, without written permission from the publisher.

John Benjamins Publishing Co. · P.O. Box 36224 · 1020 ME Amsterdam · The Netherlands
John Benjamins North America · P.O. Box 27519 · Philadelphia PA 19118-0519 · USA

Table of contents

Contributors	IX
Acknowledgements	XI
Introduction	1
<i>Vyvyan Evans and Stéphanie Pourcel</i>	
Part I. Approaches to semantics: Theory and method	
Meaning as input: The instructional perspective	15
<i>Peter Harder</i>	
Semantic representation in LCCM Theory	27
<i>Vyvyan Evans</i>	
Behavioral profiles: A corpus-based approach to cognitive semantic analysis	57
<i>Stefan Th. Gries and Dagmar Divjak</i>	
Polysemy, syntax, and variation: A usage-based method for Cognitive Semantics	77
<i>Dylan Glynn</i>	
Part II. Approaches to metaphor and blending: Theory and method	
Solving the riddle of metaphor: A salience-based model for metaphorical interpretation in a discourse context	107
<i>Mimi Ziwei Huang</i>	
When is a linguistic metaphor a conceptual metaphor?	127
<i>Daniel Casasanto</i>	
Generalized integration networks	147
<i>Gilles Fauconnier</i>	
Genitives and proper names in constructional blends	161
<i>Barbara Dancygier</i>	

Part III. Approaches to grammar: Theory and method

- What's (in) a construction? Complete inheritance vs. full-entry models 185
Arne Zeschel
- Words as constructions 201
Ewa Dąbrowska
- Constructions and constructional meaning 225
Ronald W. Langacker
- Partonomic structures in syntax 269
Edith A. Moravcsik

Part IV. Language, embodiment and cognition: Theory and application

- Language as a biocultural niche and social institution 289
Chris Sinha
- Understanding embodiment: Psychophysiological models in traditional medical systems 311
Magda Altman
- Get* and the grasp schema: A new approach to conceptual modelling in image schema semantics 331
Paul Chilton
- Motion scenarios in cognitive processes 371
Stéphanie Pourcel

Part V. Extensions and applications of cognitive linguistics

- Toward a social cognitive linguistics 395
William Croft
- Cognitive and linguistic factors in evaluating text quality: Global versus local? 421
Ruth A. Berman and Bracha Nir
- Reference points and dominions in narratives: A discourse level exploration of the reference point model of anaphora 441
Sarah van Vliet

The dream as blend in David Lynch's <i>Mulholland Drive</i> <i>Johanna Rubba</i>	465
“I was in that room!”: Conceptual integration of content and context in a writer's vs. a prosecutor's description of a murder <i>Esther Pascual</i>	499
Index	517

Contributors

Magda Altman
University of KwaZulu-Natal

Ruth Berman
University of Tel-Aviv

Daniel Casasanto
Max Planck Institute for
Psycholinguistics, Nijmegen

Paul Chilton
University of Lancaster

William Croft
University of New Mexico

Ewa Dąbrowska
University of Sheffield

Barbara Dancygier
University of British Columbia

Dagmar Divjak
Science Foundation Flanders &
University of Sheffield

Vyvyan Evans
Bangor University

Gilles Fauconnier
University of California, San Diego

Dylan Glynn
University of Lund

Stefan Th. Gries
University of California, Santa Barbara

Peter Harder
University of Copenhagen

Mimi Ziwei Huang
University of Northumbria

Ronald W. Langacker
University of California, San Diego

Edith Moravcsik
University of Wisconsin-Milwaukee

Bracha Nir
Haifa University

Esther Pascual
Groningen University

Stéphanie Pourcel
Bangor University

Johanna Rubba
California Polytechnic State University,
San Luis Obispo

Chris Sinha
University of Portsmouth

Sarah van Vliet
VU University Amsterdam

Arne Zeschel
University of Southern Denmark

Acknowledgements

For three windy days in October 2005, nearly 180 cognitive linguists descended upon Brighton, on the South coast of England, for the UK's inaugural cognitive linguistics conference. The conference, held between October 23rd and the 25th was entitled *New Directions in Cognitive Linguistics*. The make-up of delegates who attended the conference was truly international in nature, coming from all over the globe, including both North and South America, Northern and Southern Africa, the Middle East, the Far East, Australia, and from Western, Central and Eastern Europe. The conference, held on the campus of the University of Sussex, was full of excitement, intellectual curiosity and good humour, which achieved an acme, of sorts, with a final gala dinner and karaoke next to Brighton's famous Royal Pavilion. The volume you hold in your hands represents a selection of peer-reviewed papers, based, for the most part, on contributions from the conference. It captures some of the diversity and excitement of the ideas discussed and presented at the conference. In all cases, the papers contained here represent, in various ways, new directions in cognitive linguistics. Needless to say, the present volume would not have been possible without the participation and involvement of the delegates at the 2005 event. We gratefully acknowledge and thank all who attended and presented at the conference. In addition, we also wish to acknowledge the financial support of the British Academy, which, in part, made the conference possible.

We also owe a significant debt of thanks to the contributors of the present volume. A collection of this sort involves a mammoth collective effort of time, industry, and goodwill. We are extremely grateful to all the contributors for their patience and hard-work in producing such a broad range of fine contributions, and for their generosity in being willing to share the results of their research efforts. In addition, we are indebted to all the scholars who graciously assisted in the review, quality-control and selection process, as well as to our local proof-reader, Eileen Hall.

Finally, we wish to acknowledge the support of our publisher John Benjamins, particularly Seline Benjamins, Jan-Reijer Groesbeek, Kees Vaes and Hanneke Brintjes, in facilitating the publication of this volume.

Vyvyan Evans and Stéphanie Pourcel
Bangor, October 2008

Introduction

Vyvyan Evans and Stéphanie Pourcel

Since the publication of the seminal *Metaphors We Live By*, in 1980, cognitive linguistics has emerged as one of the most innovative and exciting paradigms in the interdisciplinary project known as cognitive science (see Evans and Green 2006 for a comprehensive overview of cognitive linguistics). In nearly three decades, the field has established itself at the forefront of work on grammar, linguistic semantics and aspects of conceptual structure and mental representation, to name but a few exemplars. The field also represents one of the fastest growing schools in linguistics and today boasts a sophisticated and well-established international infrastructure. In addition to a series of large-scale biennial conferences, organised under the auspices of the International Cognitive Linguistics Association (ICLA), cognitive linguistics features a significant number of national cognitive linguistics associations, affiliated to ICLA.

One symptom of the success of cognitive linguistics has been its institutionalisation, with a plethora of conferences and associations developed to enshrine its assumptions, methodologies and main theoretical paradigms. Another is the by now voluminous literature, including a detailed and sophisticated body of work in the main theoretical paradigms which populate cognitive linguistics, as well as a range of textbooks (e.g. Croft and Cruse 2004; Evans and Green 2006; Lee 2001; Ungerer and Schmid 2006), works of reference (Geeraerts and Cuyckens 2007; Evans et al. 2007), and so on, with a wide-ranging and complex technical vocabulary (see Evans 2007 for an overview).

New directions

The key objective of this 'New Directions' volume is to further contribute to this rich body of literature by firstly, taking stock of what cognitive linguistics, as an enterprise, has achieved, and secondly, by examining new avenues of investigation and exploration, new methods, new analytical means, and new ideas. The volume provides a venue for the survey of both the state of the art and new directions in cognitive linguistics. In particular, the volume surveys recent empirical and methodological trends, as well as applications of cognitive linguistics to a range of issues in neighbouring and cognate disciplines, such as psychology, sociology, anthropology, education, applied linguistics, literary studies, and more.

The papers in this volume collectively review a range of established phenomena and theories in cognitive linguistics, including approaches to figurative language, lexicalisation

patterns, meaning construction, cross-linguistic variation, grammar, and the relationship between language, conceptual structure and experience. The volume also examines and charts new directions in these areas. In addition, the volume showcases a representative selection of both the state of the art and the new in terms of methodological and empirical approaches deployed in cognitive linguistics. A further contribution made in the volume is the exploration of new areas of research, for example, cognitive sociolinguistics, and the evolutionary basis of language, as well as the exploration and presentation of recent trends in the application of cognitive linguistics to the analysis of text, narrative, discourse, dream, and film, as represented, in particular, in the final section of the book.

In essence, this volume is a testament to the wide-ranging research profile that the cognitive linguistics enterprise has developed since its inception, as well as to recent innovations. It offers both a representative sample of current practice and areas of enquiry in cognitive linguistics, as well as new trends, which seek to explore previously uncharted realms of investigation, both within the field and beyond its traditional boundaries.

An overview of the volume

The volume is divided into five sections. The first four treat traditional areas of investigation and theory in cognitive linguistics: *Approaches to semantics*, *Approaches to metaphor and blending*, *Approaches to grammar*, and *Language, embodiment and cognition*. The fifth section deals with *Extensions and applications of cognitive linguistics*. Below we preview each of the sections and the chapters contained.

I. Approaches to semantics: Theory and method

This section of the book addresses theoretical, methodological and empirical issues in cognitive semantics. The first chapter, by **Peter Harder**, *Meaning as input: The instructional perspective*, is primarily concerned with the risk of ‘usage fundamentalism’ in cognitive linguistics. This concerns the assumption that only actual utterances in fact exist. According to Harder, this position stands in opposition to the classical error of situating the truth about language at the level of abstract ideal objects. In particular, Harder is concerned as to the way in which the term ‘meaning’ is being deployed in recent work in cognitive linguistics (cf. e.g. Croft 2000; Evans 2006, this volume), and the dissociation between ‘meaning’ on one hand and ‘mental representation’ (i.e. knowledge of language) on the other. Harder argues that if meaning continues to be equated with language use (rather than knowledge of language), and this definition becomes accepted, it is no longer obvious exactly what constitutes the content side (semantic pole) of a linguistic unit. In order to remedy this, Harder presents an approach focused on a tripartition of the canonical language event into input, processing and output. The idea is that in order to choose a specific linguistic item competently, one must know what ‘input content’ it can add to the message. In order to actually succeed in making a contribution, the linguistic item has to be processed by the addressee, resulting in an understanding that constitutes the ‘output’ (as an

actual usage event). Knowing a language, he argues, consists of knowing the input properties of the forms selected by the language producer – whereas actual outputs can never be known for certain in advance. Harder outlines the basic features of such an ‘instructional’ perspective illustrating how this may serve a purpose in the current usage-oriented as well as socially-oriented trend in cognitive linguistics.

The second chapter in section I, by Vyvyan Evans, is entitled *Semantic representation in LCCM Theory*. This paper focuses on the nature of semantic representation from the perspective of the Theory of Lexical Concepts and Cognitive Models, also known as LCCM Theory (Evans 2006, To appear). LCCM Theory takes its name from the two central theoretical constructs adopted in the theory: the lexical concept and the cognitive model. The lexical concept represents the means adopted in LCCM Theory of modelling units of semantic structure. In contrast, a cognitive model is a component of conceptual knowledge, which is to say, non-linguistic knowledge. Hence, the cognitive model models units of conceptual structure. LCCM Theory assumes that lexical concepts and cognitive models are types of knowledge belonging to two distinct representational systems, which have distinct and divergent functions. These are the linguistic system, which encodes semantic structure, and the conceptual system which encodes conceptual structure. Evans argues that the linguistic system evolved, in part, by facilitating more effective control of the extant representations in the conceptual system – representations which evolved for action and perception, i.e. for non-linguistic purposes. In essence, the central argument of the paper is that the semantic representations in the linguistic and conceptual systems interact for purposes of linguistically-mediated communication. Together, the lexical concept and the cognitive model form a level of representation that the author refers to as *semantic representation*. The paper describes the nature of the lexical concept, the nature of the cognitive model, and the nature of the interaction between the two.

While the first two chapters were more theoretically-oriented, the final two chapters in part I are more concerned with method. The first of these, by Stefan Th. Gries and Dagmar Divjak is entitled: *Behavioral profiles: A corpus-based approach to cognitive semantic analysis*. One of the areas which has most strongly supported the emergence of cognitive linguistics as a new research paradigm is that of lexical semantics. Early work, in particular on prepositions, introduced the notions of prototypes, network representations and radial categories into linguistics. These innovations of cognitive-linguistic lexical semantic analysis were later used for analysing constructional elements. While this work has provided a wealth of insights, the approach – in particular the then widely used network representations of word senses – was criticised for a variety of methodological and conceptual shortcomings. It is probably fair to say that, in spite of a growing recognition of such shortcomings, cognitive linguistics is still far from having resolved all of its issues. In response, Gries and Divjak survey a variety of quantitative, corpus-based methods that can be used to pursue cognitively-inspired lexical semantic analyses. After a brief discussion of the main contributions to the field, Gries and Divjak propose quantitative techniques for addressing some of the long-standing problems in the domains of polysemy and near synonymy. In so doing, they build on previously unrelated proposals from corpus linguistics in general and corpus-based lexicography in particular. They illustrate their proposal on the basis of two case studies: the first presents selected results from a study on

the senses of a highly polysemous English verb *run*; the second applies their methodology to nine near synonymous Russian verbs meaning *try*. The semantic issues investigated in the case studies include prototype identification, the (degree of) sense distinctness, and the structure of the hypothesised network.

The fourth and final chapter in section I, by **Dylan Glynn**, is entitled: *Polysemy, syntax, and variation. A usage-based method for cognitive semantics*. In this chapter, Glynn addresses issues in the description of polysemy. He argues that results derived from the Lexical Network Model (Lakoff 1987; Cuyckens 1995) have been demonstrated to be *ad hoc* (Sandra and Rice 1995; Tyler and Evans 2001). He suggests that while the Principled Polysemy framework (Evans 2005) improves on this model with a more constrained analytical apparatus, a radically different yet complementary model is, nevertheless, required. Accordingly, Glynn presents a usage-based quantitative and multifactorial method that adheres to the theoretical tenets of cognitive linguistics (Langacker 1987; Lakoff 1987) and draws from existing methodologies in the study of near-synonymy (Geeraerts et al. 1994; Fischer 2000; and Gries 2003). The method uses feature analysis of different variables and employs correspondence analysis to reveal feature association. Glynn argues that the resulting clusters of features represent polysemic structure.

In sum, the four papers in this section represent an overview of some of the recent theoretical controversies in the arena of cognitive approaches to semantics, and new directions, both theoretical and methodological, which attempt to resolve some of these outstanding issues.

II. Approaches to metaphor and blending: Theory and method

Section II of the book is concerned with the two phenomena known as metaphor, and variously conceptual integration or blending. The first two chapters deal with metaphor, while the second two are concerned with blending. The chapters collectively address both theoretical and methodological issues, as well as examine these phenomena in new ways and contexts,

The first chapter, by Mimi Ziwei Huang, is entitled: *Solving the riddle of metaphor: A salience-based model for metaphorical interpretation in a discourse context*. The purpose of this chapter is to examine how metaphor is interpreted in a discourse context. Huang employs the Graded Salience Hypothesis (Giora 1997) in order to do so. She argues that three salient factors are decisive in metaphorical interpretation. The first is the graded salient lexical meaning of a word or an expression, together with its semantic fields and scenarios. The second is the metaphorical mapping process contributed to by the metaphorical source, target, co-text and context. The third salient factor is the intended metaphorical meaning in a given context. Huang illustrates these three salient factors by virtue of an analytical account of a short story taken from *The Devil's Larder* (Crace 2002).

The second chapter in this section, by **Daniel Casasanto**, is entitled: *When is a linguistic metaphor a conceptual metaphor?* In his chapter, Casasanto is concerned with establishing whether conceptual metaphors have psychological reality. According to Conceptual Metaphor Theory, metaphors are fundamentally *conceptual* structures – not linguistic structures (Lakoff 1993). Yet, the majority of evidence for conceptual metaphors comes

from analysis of linguistic metaphors. Casasanto asks whether we can necessarily infer how people think from the way they talk. This chapter illustrates some dangers of building a theory of concepts principally upon linguistic data. The chapter briefly reviews experimental work testing our understanding of the abstract domain of *time*, and then presents experiments testing the metaphorical basis of *similarity*. Three experiments tested the relationship between similarity and spatial proximity predicted by Conceptual Metaphor Theory (Lakoff and Johnson 1980, 1999). In all experiments similarity ratings for pairs of words or pictures varied as a function of how far apart stimuli appeared on the computer screen, but the effect of distance on similarity differed depending on the type of judgments participants made. Stimuli presented closer together were rated more similar during 'conceptual' judgments of abstract entities or unseen object properties, consistent with predictions based on linguistic metaphors. By contrast, stimuli presented closer together were rated *less* similar during 'perceptual' judgments of visual appearance, contrary to the conceptual metaphor SIMILARITY IS PROXIMITY. Casasanto argues that these results underscore the importance of testing Conceptual Metaphor Theory experimentally, and suggest that linguistic metaphors should be treated as a source of *hypotheses* about conceptual structure – hypotheses that require both linguistic and extra-linguistic evaluation.

The third chapter in section II, by Gilles Fauconnier, is entitled: *Generalized Integration Networks*, and deals with Blending or conceptual integration. Fauconnier argues that the systematic study of integration as a cognitive operation made many useful descriptive distinctions possible. So, within the data referred to as "blends", there are many different products depending on the types of inputs, the links between them, the choices for projection, and so forth. Corresponding types of blends have been distinguished, or rather aligned on a graded continuum, going from simplex blends to mirror blends to single-scope and double-scope blends, all dividable into further subcategories (Fauconnier and Turner 2002). While, according to Fauconnier, the description and classification of this new data is largely uncontroversial and widely viewed as innovative and useful, a more significant project is to explore the role of integration and compression in meaning construction beyond these very visible blends. Accordingly, in his chapter, Fauconnier, points out some useful generalisations that emerge from the study of integration, along with some of the pervasive fallacies that stand in the way of making such generalisations. Through the analysis of attested data, he discusses the notion of "generalized integration networks" and how they allow the construction of a multiplicity of surface products in human thought and action.

Like the chapter by Fauconnier, the fourth and final chapter in section II also addresses blending. The contribution by Barbara Dancygier entitled: *Genitives and proper names in constructional blends*, presents a blending analysis of genitives, thereby providing a theoretical and methodological illustration of the role of blending in language as well as illustrating the utility of conceptual integration as a theoretical construct. According to Dancygier, the genitive (’s) form in English has long been seen as semantically puzzling. It plays a special role as the only case in English which is morphologically marked on nouns, and displays a very broad array of meanings and uses (Nikiforidou 1991; Taylor 1996; Rosenbach 2002). The recent view of the genitive is that it is a means of establishing a reference point (Langacker 1991; Taylor 1996) for the construct represented by the

noun being modified. In her chapter, Dancygier describes a somewhat more specific use of the genitive, which emerges as the specific contribution of the genitive to two syntactic constructions, both of which can be represented as conceptual integration networks. Dancygier discusses the contribution of the genitive to constructional meaning in terms of two theoretical constructs: *frame metonymy* and *constructional compositionality*. These notions are illustrated by virtue of an analysis of two constructions. The first Dancygier terms the *GEN-XYZ construction*, exemplified by the sentence *Too much of the world, Cambodia has become "Vietnam's Vietnam."* The second construction which relies on a similar use of the genitive is *One person's X is another person's Y*, represented in *One person's trash is another person's treasure*.

III. Approaches to grammar: Theory and method

Section III of the book is concerned with cognitive linguistic approaches to grammar, as in previous sections addressing and assessing recent trends and perspectives, both theoretically and methodologically, and charting new issues and avenues for exploration. The section opens with the chapter from Arne Zeschel entitled: *What's (in) a construction? Complete inheritance vs. full-entry models*. Zeschel's chapter contrasts the two most widely assumed criteria for constructional status that have been proposed in the Construction Grammar literature. Departing from a corpus study of a particular 'schematic idiom' of English, the chapter presents both theoretical and empirical arguments for a usage-based interpretation of the term *grammatical construction* that accords unit status to linguistic elements that are sufficiently *entrenched*. Zeschel argues that the criterion of non-predictability that is often employed in computational approaches is inappropriate for accommodating the inherently flexible and creative aspects of human problem solving that are exhibited by naturally occurring language.

The second chapter in section III also takes up the issue of the nature and status of a construction. In her chapter entitled: *Words as Constructions*, Ewa Dąbrowska explores the status of words as linguistic units. She argues that the lexical representations of verbs and other relational predications include schematic specifications of the entities participating in the relationship as well as salient aspects of the setting and can thus be regarded as constructions (Langacker 1987, 2005). From this perspective, a unified account is afforded of how lexical knowledge is acquired and represented. Moreover, this also facilitates understanding how, in the later stages of language development, learners are able to construct detailed lexical entries for verbs by relying mainly or exclusively on information about typical collocational patterns. Her chapter provides further empirical support for this view by describing two experiments tapping adult speakers' knowledge about 18 relatively low-frequency verbs designating manner of walking and running (*trudge, plod, scurry, scamper*, and so on.)

The third chapter in section III, by Ronald W. Langacker, is entitled: *Constructions and constructional meaning*. Langacker outlines a unified approach to a range of language-related phenomena involving conceptual structures linked by directional relationships, whereby one structure is invoked to apprehend another. Among these phenomena are inheritance, categorisation, composition, derivation, metaphor, and blending. According to Langacker,

from this perspective, there is no clear distinction between syntagmatic and paradigmatic relationships, as constructional schemas contribute to a complex expression's meaning in the same way that component structures do. A lexeme's grammatical category is indissociable from the constructions it appears in, which in turn are part of its characterisation. Lexical and constructional meaning overlap and are often indistinguishable.

The chapter also addresses, in this context, the issue of whether lexemes develop meanings appropriate to the constructions they occur in, or whether construction-specific aspects of meaning remain the province of the constructions themselves.

The final chapter in this section, by **Edith Moravcsik**, is entitled: *Partonomic structures in syntax*. While the preceding three chapters dealt with the nature and status of constructions and constructional meaning, Moravcsik addresses the related issue of part-whole organisation in grammatical structure. Her chapter seeks to demonstrate two ways in which partonomy – whole-part relations – is a useful conceptual tool in formulating generalisations in syntax. First, she claims that positing phrases and clauses as wholes that subsume words as their parts simplifies the statement of syntactic rules. However, syntactic wholes may be problematic: they are often complex and evidence for them is frequently contradictory. According to Moravcsik, in several syntactic frameworks, such complexities and inconsistencies are resolved by slicing a single structure into layers – a second application of partonomic analysis. Moravcsik concludes by providing examples of similar uses of partonomy outside linguistics, which, she suggests, highlight a common cognitive component of argumentation in linguistics, science, and everyday thinking.

IV. Language, embodiment and cognition: Theory and application

This section of the book addresses an important strand of research in cognitive linguistics, namely the relationship between language, embodiment and cognitive structure and function. The four chapters in this section of the book address new perspectives on several aspects of this interface.

The first chapter, by **Chris Sinha**, entitled: *Language as a biocultural niche and social institution*, outlines a biocultural theory of language and its acquisition. In so doing, Chris Sinha examines the relationship between the emergence of language, culture and action. Moreover, he situates the emergence of mental grammar in an evolutionary context, viewing language as the outcome of the more general development of human semiotic abilities. From the perspective of the biocultural theory that he develops, the view of grammar that emerges is, in one sense richer, and in another poorer, than that to which we have become accustomed from outside cognitive linguistics. It is richer because it incorporates meaning and context, the twin pillars supporting both language acquisition and language use. It is poorer because there is no longer a compelling reason to attribute a knowledge equivalent to the results of formal analysis to the learners and users of language. Hence, there is no mental grammar isomorphic with autonomous grammar, as in generative linguistics, for instance. Rather, grammar is *in language*, as a biocultural niche and social institution. According to Sinha, the learner need not internalise a formal description of the structure of language in order to acquire the ability to *act* in it. Language is not an “input” to a processor or device, but a structured niche affording complex and semiotically mediated

communication and cognition. On this view, grammar is a social institution, regulating linguistic practice, and it is the practical ability to adhere to the constraints and supports imposed by and related to language that is acquired by the language learner.

The second chapter in section IV, by **Magda Altman**, is entitled: *Understanding embodiment: Psychophysiological models in traditional medical systems*. While the notion of embodiment has been highly influential in cognitive linguistics and, more generally, cognitive science, Altman argues that representations of the body in traditional medical systems are a largely neglected yet invaluable resource for information on pretheoretical conceptualisations of the body. Accordingly, the chapter analyses several traditional Chinese psychophysiological models with reference to current work on mimesis, image schemas, perception, proprioception and the body schema. Altman's investigation suggests that traditional medical systems may capture the subjective experience of embodiment in a structured and integrated manner complementing other methods of scientific investigation in helping us understand the nature of embodiment and the kind of 'body' that language instantiates.

The third chapter, by **Paul Chilton**, is entitled: *GET and the GRASP schema: A new approach to conceptual modelling in image schema semantics*. Chilton's starting point is the observation that *get* is a verb with a wide range of uses that are not obviously related to one another. His chapter represents an attempt to make progress toward a unified account by introducing a novel theoretical framework. This framework relies heavily on spatial concepts formalised in informal geometric terms; crucially, the framework integrates foregrounding/backgrounding in discourse, temporal viewpoint, modal distance and directionality (Chilton 2005, 2007). The fundamental principles of the approach are cognitive: it is proposed that the construction meanings associated with *get* are a conceptual category revolving around a prototype whose meaning is embodied in an image schema.

The final chapter in section IV is by **Stéphanie Pourcel**. This chapter is entitled *Motion scenarios in cognitive processes* and addresses the issue of linguistic relativity, an issue which is becoming of increasing importance in cognitive linguistics. Pourcel tests linguistic relativity (i.e. the effects of language forms on cognition) by examining memory and inference in the conceptualisation of motion. Most previous work on linguistic relativity has focused on testing the cognitive functions of categorisation and memory. Few studies have found any effects and possibly none have explored inferencing. Inferencing is an extremely promising avenue for investigating the potential influences of language on cognition, as it relies heavily on other cognitive processes, such as attention and memory. Accordingly, it potentially offers an insightful index of the relativity of conceptualisation via these other cognitive modes of processing, and hence constitutes a new direction in relativistic research. Tests were performed with English and French native speakers to offer a comparative assessment of the potential relativity of inferencing motion event information, based on the same objective stimulus. Results present considerable differences across the two language groups. English native-speaker performance reveals more manner-salient conceptualisation, while French native-speaker performance reveals more path-salient conceptualisation. In essence, Pourcel's chapter offers innovative insights in terms of new experimental methods for investigating the relativistic influence of language on non-linguistic cognition.

V. Extensions and applications of cognitive linguistics

In this section of the book, contributions focus on either extending cognitive linguistics beyond its tradition purview, as in the first chapter in this section, or in applying cognitive linguistics to text, narrative, film and judicial discourse analysis, as in the remaining chapters.

The first chapter in this section, by **William Croft**, is entitled: *Toward a social cognitive linguistics*. Croft's premise in his chapter is that in order to be successful, cognitive linguistics must go 'outside the head' and incorporate a social-interactional perspective on the nature of language. Croft sets himself the task, in this chapter, of doing exactly this. He does so by attempting to integrate foundational work in pragmatics and sociolinguistics. In particular, Croft draws on the interpretation of the pragmatic research by the psycholinguist Herbert H. Clark, who has argued for a comprehensively and consistently social cognitive perspective on language (e.g. Clark 1996). He also takes inspiration from the work of the psychologist Michael Tomasello (e.g. Tomasello 1999, 2003). Once a synthesis between these traditions has been developed, Croft demonstrates what he takes to be the fruitfulness of this approach by addressing traditional cognitive linguistic questions, in particular the nature of construal and its relation to grammar.

The next chapter in the final section, by **Ruth Berman and Bracha Nir**, is entitled: *Cognitive and linguistic factors in evaluating text quality: Global versus local?* This chapter seeks to shed light on the elusive notion of 'text quality' across later language development. Specifically, Berman and Nir consider text construction abilities in terms of the interplay between cognition and language in the integration of bottom-up and top-down cognitive processes and the question of whether and how these relate to local linguistic expression as compared with global discourse organisation. To this end, the chapter reports on a study which examined the connection between quantitative and qualitative text measures applied to essays written by 160 children, adolescents, and adults. The texts were analysed along two distinct dimensions: (1) local linguistic expression in lexical usage (by such measures as word length, lexical density, and linguistic register) and syntactic constructions (e.g. clause length, noun phrase complexity, relative clauses) and (2) global discourse quality – by means of an innovative analytical framework. The results reveal a marked correlation between different measures of lexical usage, on the one hand, and between these and the use of syntactic constructions, on the other. Yet a dissociation emerged between measures of local language use and of global text quality. These findings are discussed as reflecting corresponding but not necessarily interdependent patterns of linguistic and cognitive development across adolescence and as suggesting caution in equating complicated, high-level linguistic expression with more general text construction abilities.

The next chapter, by **Sarah van Vliet**, is entitled: *Reference points and dominions in narratives: A discourse level exploration of the reference point model of anaphora*. In this chapter, van Vliet presents a discourse level exploration of Van Hoek's (1997) *reference point model of anaphora*. Within this model, the felicitous use of coreferential full nominals and pronouns depends on the extent to which a nominal is construed as a conceptual reference point within the immediate context. Van Vliet's chapter aims to demonstrate that this characterisation may also account for the alternate use of proper

nouns and pronouns in narratives, given a sufficiently detailed account of discourse context. The chapter describes a number of context factors – such as episode structure, referential distance, point of view and competing referents – which exert their influence on referential form as part of attention framing throughout discourse.

The fourth chapter in section V, by **Johanna Rubba**, is entitled: *The Dream as blend in David Lynch's Mulholland Drive*. This chapter applies Conceptual Blending Theory (Fauconnier and Turner 2002) to David Lynch's film *Mulholland Drive*. Rubba's interpretation of the film's plot is that the first two-thirds of its running time correspond to the protagonist's dream, an attempt to repair a life that has gone horrifically wrong. According to Rubba, dreams can be seen as self-contained mini-worlds of conceptual projections from our experience of waking life. Yet, while the stuff of dreams is taken from our conceptual structure, dreams have their own logic, revising, distorting, and defying reality. This makes them an interesting object for study within the Conceptual Blending framework. Rubba's study makes a compelling case for dreams as blends, showing how Blending Theory also provides tools for other correspondences between dream worlds and waking worlds. In essence, Rubba's chapter presents a fascinating case study of the application of one particularly well-known cognitive linguistic theory to the analysis of the cinema of David Lynch.

The final chapter in this section is by **Esther Pascual**. The paper is entitled "I was in that room!": *Conceptual integration of content and context in a writer's vs. a prosecutor's description of a murder*. Pascual's chapter comprises a case study in the discourse concerning a murder. Like the previous chapter, Pascual applies Conceptual Blending Theory in conducting her analysis. She proceeds by comparing a writer's and a prosecutor's description of a murder, both involving the integration of the discourse content with the communicative context. Pascual argues that this content-context blend is essential to the meaning and communicative effect of these discourses. This blend also seems fundamental to inter-subjectivity skills such as empathy, ordinary understanding, and the use of the generic 'you'.

References

- Chilton, P. 2005. Discourse Space Theory: Geometry, brain and shifting viewpoints. *Annual Review of Cognitive Linguistics* 3: 78–116.
- Chilton, P. 2007. Geometrical concepts at the interface of formal and cognitive models: *Aktionsart*, aspect and the English progressive. *Pragmatics and Cognition* 15 (1): 91–114.
- Clark, H. H. 1996. *Using Language*. Cambridge: Cambridge University Press.
- Crace, J. 2002. *The Devil's Larder*. London: Penguin.
- Croft, W. 2000. *Explaining Language Change*. London: Longman.
- Croft, W. & Cruse, D. A. 2004. *Cognitive Linguistics*. Cambridge: Cambridge University Press.
- Cuyckens, H. 1995. Family resemblance in the Dutch spatial prepositions *Door* and *Langs*. *Cognitive Linguistics* 6: 183–207.
- Evans, V. 2005. The meaning of *Time*: Polysemy, the lexicon and conceptual structure. *Journal of Linguistics* 41 (1): 33–75.
- Evans, V. 2006. Lexical concepts, cognitive models and meaning-construction. *Cognitive Linguistics* 17 (4): 491–534.
- Evans, V. 2007. *A Glossary of Cognitive Linguistics*. Edinburgh: Edinburgh University Press.

- Evans, V. This volume. *Semantic representation in LCCM Theory*.
- Evans, V. To appear. *How Words Mean*. Oxford: Oxford University Press.
- Evans, V., Bergen, B. K. & Zinken, J. (eds.). 2007. *The Cognitive Linguistics Reader*. London: Equinox.
- Evans, V. & Green, M. 2006. *Cognitive Linguistics: An Introduction*. Edinburgh: Edinburgh University Press.
- Fischer, K. 2000. *From Cognitive Semantics to Lexical Pragmatics. The Functional Polysemy of Discourse Particles*. Berlin: Mouton de Gruyter.
- Fauconnier, G. & Turner, M. 2002. *The Way We Think*. New York: Basic Books.
- Geraerts D. & Cuyckens, H. (eds). 2007. *Handbook of Cognitive Linguistics*. Oxford: Oxford University Press.
- Geraerts D., Grondelaers, S. & Bakema, P. 1994. *The Structure of Lexical Variation. Meaning, Naming and Context*. Berlin: Mouton de Gruyter.
- Giora, R. 1997. Understanding figurative and literal language: the Graded Salience Hypothesis. *Cognitive Linguistics 7*: 183–206.
- Gries, S. Th. 2003. *Multifactorial Analysis in Corpus Linguistics: A Study of Particle Placement*. London: Continuum.
- Lakoff, G. 1987. *Women, Fire and Dangerous Things*. Chicago: University of Chicago Press.
- Lakoff, G. 1993. The contemporary theory of metaphor. In *Metaphor and Thought*, A. Ortony (ed.), 202–251. Cambridge: Cambridge University Press.
- Lakoff, G. & Johnson, M. 1980. *Metaphors We Live By*. Chicago: University of Chicago Press.
- Lakoff, G. & Johnson, M. 1999. *Philosophy in the Flesh: The Embodied Mind and its Challenge to Western Thought*. Chicago: University of Chicago Press.
- Langacker, R. W. 1987. *Foundations of Cognitive Grammar. Theoretical Prerequisites*. Volume 1. Stanford: Stanford University Press.
- Langacker, R. W. 1991. *Foundations of Cognitive Grammar. Descriptive Applications*. Volume 2. Stanford: Stanford University Press.
- Langacker, R. W. 2005. Construction Grammars: cognitive, radical, and less so. In *Cognitive Linguistics. Internal Dynamics and Interdisciplinary Interaction*, F. J. Ruiz de Mendoza Ibanez & M. S. Pena Cervel (eds), 101–159. Berlin: Mouton de Gruyter.
- Lee, D. 2001. *Cognitive Linguistics: An Introduction*. Oxford: Oxford University Press.
- Nikiforidou, K. 1991. The meanings of the genitive: A case study in semantic structure and semantic change. *Cognitive Linguistics 2* (2): 149–206.
- Rosenbach, A. 2002. *Genitive Variation in English: Conceptual Factors in Synchronic and Diachronic Studies*. Berlin: Mouton de Gruyter.
- Sandra, D. & Rice, S. 1995. Network analyses of prepositional meaning: Mirroring whose mind – the linguist's or the language user's? *Cognitive Linguistics 6*: 89–130.
- Taylor, J. R. 1996. On running and jogging. *Cognitive Linguistics 7* (1): 21–34.
- Tyler, A. & Evans, V. 2001. Reconsidering prepositional polysemy networks: The case of *over*. *Language 77* (4): 724–765.
- Tomasello, M. 1999. *The Cultural Origins of Human Cognition*. Harvard: Harvard University Press.
- Tomasello, M. 2003. *Constructing a Language*. Harvard: Harvard University Press.
- Ungerer, F. & Schmid, H. J. 2006. *An Introduction to Cognitive Linguistics*. London: Longman.
- Van Hoek, K. 1997. *Anaphora and Conceptual Structure*. Chicago: The University of Chicago Press.

PART I

Approaches to semantics

Theory and method

Meaning as input

The instructional perspective

Peter Harder

1. Introduction

The widespread trend to expand into the social sphere as part of constructing a comprehensive usage-based linguistics is one of the most promising developments in current cognitive linguistics. Bringing together cognitive structures and interactive practices, and providing an architecture in which actual language use is the foundation, this approach seems well placed to avoid the traditional sources of error in the linguistic tradition. These centrally include the practice of looking for timeless, abstract entities as the locus of linguistic reality. In the case of Saussurean structuralism, the essence of language is a socially shared immanent 'langue' that subsumes all varieties under one overarching formula; in the case of Chomskyan generative grammar, the truth about language resides in the internal grammar that stands above the vagaries of E-language. Like defunct prescriptive grammar, both reduce actual usage to the poor country cousin of 'core' linguistic facts.

No generally accepted format for tackling the whole process of language use has yet appeared, however. This article is an attempt to highlight one necessary element in establishing a complete strategy of usage-oriented realism, namely the role of linguistic items in starting off, or triggering, processes of understanding – as opposed to their role in standing for completed instances of utterance understanding. I use the term 'instructional semantics' about the dimension I want to emphasize, although the word 'instruction' may ring overly 'imperative': the point is that words can usefully be understood as designed to prod, or prompt, the addressee to carry out interpretive activities of specifiable kinds. This angle has been brought up from time to time (cf. vol. 39, a special issue of *Acta Linguistica Hafniensia*), but has yet to seriously capture the imagination of linguists. I believe it has a key role to play in the more dynamic and user-oriented conception of linguistics that is now emerging. Among cognitive linguists who include this angle, Fauconnier (1985: 2) uses the word 'instruction' about the "underspecified" meanings of linguistic items in calling upon the fully specified cognitive representations. At that time Fauconnier rightly argued that it was necessary to move towards the full cognitive perspective rather than limit one's attention to the purely linguistic input; I argue that more than twenty years later it is time to go back to the specifically linguistic perspective in order to be precise about its role in the full picture. In doing so I align myself with the position on polysemy taken by Tyler and Evans (2003: 40), who use the word 'prompts' about the contribution

of words to utterance meaning, and with Evans (2006) in distinguishing lexical concepts from conceptual models.

The motivation for bringing it to attention as part of the new developments in cognitive linguistics is what I see as a risk of 'usage fundamentalism', i.e. the assumption that only actualized utterances really exist. This position constitutes the opposite extreme in relation to the classical error of locating the truth about language at the level of abstract ideal objects. Even though no one to my knowledge explicitly advocates an uncompromising version of this belief (the position of Thompson (2002) is perhaps the closest you can get; cf. the discussion in Boye and Harder 2007), I think it makes sense to discuss it as a risk that is part of the landscape when there is a trend away from abstract entities and towards concrete situated utterances.

In the context of cognitive linguistics, one question that raises the issue is what is happening to the term 'meaning', arguably the most basic notion in language and linguistics. At present the term is often used in a way that implies that meaning only exists in actual utterances (cf. e.g. Croft 2000 and Evans 2006). If this usage becomes generally accepted and combined with the well-established cognitive-linguistic position that meaning is encyclopaedic so that no specifically 'linguistic' meaning can be factored out from the whole set of cognitive and conceptual representations, it is not immediately obvious exactly what constitutes the content side (semantic pole) of a linguistic sign. Since the description of language, rather than a description of the entire cognitive universe, is what linguists do, clarity on this point is essential. Strategically, it would be a mistake to leave the question to those linguists, including generativists, who provide clear but misguided answers by positing unwarranted forms of separation between 'language' and 'actual usage'.

I would like to highlight an approach focused on a tripartition of the canonical language event into input, processing and output. (The point is most easily understood from a reception perspective, but it also applies to the production perspective: to be successful, a speaker must encode messages based on what words can do for him/her.) The idea is that in order to choose a specific linguistic item competently, one must know what 'input content' it can add to the message. In order to actually succeed in making a contribution, it has to be processed by the addressee, resulting in an understanding that constitutes the 'output' (as an actual usage event). Knowing a language, I suggest, essentially consists of knowing the input properties of the forms you choose – whereas actual outputs can never be known for certain in advance.

Below I advocate, therefore, that a complete usage-based theory should recognise the importance of specifying sign content at the stage of 'input' properties. Obviously, such an account cannot stand alone; it needs to be linked up with the other two phases of the tripartition, i.e. accounts of the type of things that happen to semantic inputs in the course of processing, as well as accounts of the properties of the 'output stage', i.e. fully specified utterance meanings. Below I try to outline the basic features of such an approach and show how this endeavour may serve a purpose in the current usage-oriented as well as socially oriented trend in cognitive linguistics.

2. The basic idea: Sources and current relevance

One of the sources of instructional thinking about meaning is the computational analogy. 'Utterances as programs' (Davies and Isard 1972) explores the parallel between understanding utterances and compiling programs viewed as structured directions for computational operations. The 'procedural-declarative' discussion, however, lost momentum after a while, because in the computational perspective it may be viewed as a matter of notational variants: an effective algorithmic procedure specifies its end state unambiguously, so you end up in the same place whether you take a procedural or a declarative view. Another source is the French linguistic tradition focusing on 'énonciation' (cf. Ducrot 1972); a recent example of how the instructional angle has been developed in that context is in the theory of 'polyphony', cf. Nølke et al. (2004). Relevance theory (cf. Sperber and Wilson 1986) has also used the idea of meaning as input to the construction of propositions rather than as coding propositions directly; the Columbia school (cf. e.g. Contini-Morava et al. 2004) also regards meaning as 'hints' rather than fully-fledged messages. In cognitive linguistics, the idea plays an important role in mental space theory, where for instance the notion of a 'space builder' exemplifies the instructional perspective.

In spite of all these converging efforts, actual convergence still remains to be established when it comes to the significance of this perspective for understanding the precise nature of coded meaning, as opposed to utterance meaning. In general, what happens is that the instructional perspective is lost from view once it has ushered in the central construct. In relevance theory, the key element is the 'explicature' that is constructed on the basis of the input. In mental space theory, the main interest is on the output in the form of the appropriate mental space configuration. A good ('effective') computer program is one that gets the right output.

But does not exactly the same thing apply to language – that the success criterion as well as the central fact about language is successful utterance understanding? Yes, but the perspective of the linguist is by definition not the same as the perspective of the language user. Language is *designed* to be overlooked, as the Danish linguist Hjelmslev said (1943:6–7), because it is the means rather than the end. Therefore language users are well advised to treat linguistic expressions as a forgettable step on the way and focus on understanding utterances. Linguists, however, cannot do their job if they follow the same advice and forget about language because actual understanding is more real.

Cognitive linguistics is under a special risk because it is in the position of challenging a mainstream approach that focuses on language by creating an unwarranted chasm between language 'itself' and language use. The basic insight that originally motivated cognitive linguistics was the realization that you cannot understand language except in the context of cognition as a whole, and to focus on 'specifically linguistic' contributions to cognition and communication therefore goes against the grain. Invoking a tripartition between input, process and product, however, is an entirely different enterprise than an artificial isolation of language from the rest of cognition and communication: if language is designed to be input to a process whose success criterion is bound up with the output end, it would obviously make no sense to study the input end in isolation. Precisely because language itself is not the end, but the means to an end, we need to allow for a perspective that factors

out – to the extent it is empirically possible – the description of the linguistic input from the description of that full cognitive-communicative understanding which constitutes the point of language use.

An example of how such a strategy may promote an integrated understanding is in the understanding of definiteness. Definiteness as a linguistic content element (the semantic pole of the definite article) can be understood as an instruction to identify the relevant discourse entity being referred to (Harder and Kock 1976; Heim 1983). Consider the following example:

After meeting in emergency session on Monday, the UN Security Council unanimously condemned the test and began negotiations on imposing tougher sanctions against Kim Jong-il's reclusive state. (Guardian Weekly, 10 October 2006)

In processing the expression *the test*, the well informed reader will insert the event of North Korea testing its first nuclear bomb as the intended referent, and this will be a crucial constituent of the understanding constructed. However, this is clearly not part of the encoded content of the expression *the test*. The linguistic expression specifies a category ('test'), and assigns definiteness to it, thus instructing the reader to identify a discourse entity belonging in that category as the intended referent. Unless the reader can do this, successful *utterance* understanding cannot come about.

However, even an uninformed reader can get at a perfectly appropriate *linguistic* understanding – and diagnose the communication problem accordingly ('I know what the phrase *the test* means, but I don't know what they're talking about'). The traditional descriptions of definiteness specify conditions that must obtain in order for the result to be obtained (existence and uniqueness, cf. Russell 1905); in the cognitive context, the role of 'mental contact' is central (cf. Langacker 1991: 98). Recognizing the status of reference as an *act* (cf. Searle 1969) is in harmony with the instructional view, but does not factor out the specific role of definiteness as the linguistic, conventional 'request' for identifying the referent (cf. Harder and Kock 1976: 24).

This is where an instructional description can enable a clear analysis by carving the tripartite process of engendering meaning-in-use at the joints: in all successful cases of language use, linguistic items must bring their encoded input to bear on cognitive and situational elements, triggering a process of understanding that succeeds in integrating all relevant linguistic and non-linguistic constituents into the utterance understanding that constitutes the final output.

Definiteness is now widely recognized as a procedural element; but the 'input' perspective is relevant for all linguistic meanings, including lexical meaning. This may not be obvious, because lexical meaning is central also in 'take home' output meaning. But as argued in Evans (2006), it is necessary to reassert the distinction between lexical (i.e. linguistic) concepts serving as prompts and cognitive models, precisely because competent use of lexical words depends on knowing specifically what the *word*, as opposed to language-independent underlying conceptual *models*, will do for you. A cognitively realistic account of how utterance meaning arises depends on being able to distinguish between lexical concepts (capturing what the speaker's knowledge of *words* brings to the process)

and the rest of the interpretive process, including the way words interface with conceptual knowledge in general.

The most salient difference is perhaps the social dimension of words as opposed to cognitive networks in the brain: only Humpty Dumpty can confidently expect words to evoke exactly what he has in mind – the rest of us need to be sure what the conventional input value of the word is in the context we are using it (cf. also Clark 1998 on ‘communal lexicons’): the issue of whether a word is appropriate is distinct from the issue of what cognitive models the speaker has in mind. Successful lexical encoding depends on using *words* in ways that will achieve intended communicative effects, not on the sum total of available conceptual models.

The three-stage approach applies to lexical meaning in the following way. At the ‘input’ stage, a lexical word functions as an instruction for the addressee to ‘access’ (Langacker’s term) the whole potential network associated with the word in order to make sense of the whole utterance. The notion of a whole potential as constituting word meaning has been put forward in various contexts, cf. Harder and Togeby (1993), Allwood (2003), Evans (2006); the differences matter less than the shared point that there are major advantages in making a basic distinction between actual and potential meaning, rather than trying to discuss homonymy, polysemy and compositionality based on actual meanings only. One of them is that the role of contextual modulation of meaning can be discussed more realistically if it is assumed that actual meanings are not enumerable in advance, but depend on what happens to the potential in actual processing situations.

3. Stability and variation

One of the chief obstacles to developing a usage-based linguistics has been the persistent Platonic heritage according to which the foundation of all knowledge must be sought in an underlying object hidden from direct inspection, based on which actual instances can be explained. Cognitive linguistics has not managed entirely to avoid getting caught on the horns of the dilemma, since it took over and transformed the Chomskyan notion of inaccessible underlying ‘tacit knowledge’, an intellectual debt recognized in Lakoff and Johnson (1999: 472), where Chomsky is said to ‘deserve enormous credit’ for this idea. The chief constituent of the stable underlying landscape is the network of mappings that link up cognitive models, including metaphorical links between domains, yielding a backbone of unconscious but stable conceptual mappings that constitute the essential reality of language. In turning away from arbitrary linguistic categories towards categories motivated by functional and cognitive factors, a natural orientation in the enterprise of cognitive linguistics had been towards stable motivational features, in the light of which linguistic features that had been understood in terms of arbitrary structure could fall into place in an underlying order that had previously been ignored.

As underlined by Croft (2000, 2001), however, the whole landscape of linguistics and language use is situated in a context of ever-active mechanisms of change, crucially including Darwinian mechanisms of selection and propagation. The previous orientation in functional and cognitive linguistics towards linguistic universals that could be naturally

associated with universals of a functional and cognitive nature has turned out to lack solid empirical foundation (cf. also the preface to Engberg-Pedersen et al. 1996). Linguistic categories cannot be assumed to be generalizable beyond the constructional contexts in which they occur. What is shared between constructions of different kinds inside one language, as well as between different languages, is not linguistic categories, but only the overall landscape with all its variations. As argued in Croft and Poole (2008), valid generalizations arise only against the background of the whole spectrum of variation, as constraints on variability rather than as underlying sameness. No forms of preconceived stability can be assumed to be underlyingly valid for any particular case.

This notion of stability as constraints on variation, I argue, is essential also in understanding the nature of linguistic meaning. From the instructional perspective, in using a particular linguistic sign, a speaker is not invoking a timeless essence but trying to reduce the spectrum of variation which the addressee's attempt to understand him/her would otherwise be subject to. The search for the type of meaning that is associated with a linguistic expression in advance of an actual occasion of use, i.e. at the input stage, should not be defined as a search for something precise and immutable, but as something that constrains variation. Here, too, you can only get at the appropriate generalization if you presuppose variability as the background against which you must operate, rather than trying to get behind it to something timeless and pristine.

Viewing meaning as process input naturally captures this kind of variation-based relative stability: the addressee's attempt to understand takes place in the midst of the flux of ongoing cognitive and social processes, but unless the speaker and addressee manage to converge to some extent, understanding is not possible. Linguistic meanings, as available to speakers in advance of a potential utterance, must be understood as means to promote the convergence. To describe them as such is therefore an essential part of the linguistic enterprise.

4. Indeterminacy: The input and the output perspective

It remains to show what the instructional dimension has to offer to the picture. The key idea is that to be a competent speaker one needs to extract something for future use from actual usage events. In order to be useful, what one extracts must be constrained in ways that converge among members of the speech community, and represent a reduction of complexity to a manageable format.

This constitutes a complementary perspective compared to the approach that locates meaning solely in actual usage events. The following quotation from Croft and Cruse (2004) can serve as an illustration of the niche I see for the instructional perspective:

In many approaches to meaning, there is a determinate starting point for the process of constructing an interpretation, but an indeterminate end point.... The present model of comprehension has an indeterminate starting point (a purport) and a determinate end point. (...).

(...) Each lexical item (word form) is associated with a body of conceptual content that is here given the name **purport** ... purport is continually developing: every experience of the use of a word modifies the word's purport to some degree. (ibid.: 100)

It is by a series of processes of construal that an essentially non-semantic purport is transformed into fully contextualized meanings... (ibid.: 103)

While there is an interpretation under which I fully agree with what is claimed here, I think the complete picture of meaning as a property of language cannot be one in which the content side of a word as part of the speaker's knowledge of language is viewed as indeterminate and "essentially non-semantic."

I would like first to address the issue of indeterminacy. Word meaning is indeterminate outside of an actual utterance in the sense that the whole potential may contain alternative options that must be resolved, and we cannot know in advance what is relevant to a particular utterance (as in the case of *bug* explored by Swinney 1979). Only when you have interpreted the word in context do you know precisely 'what it means' in that context. But part of that observation boils down to the truism that you cannot understand utterance meaning in the absence of an utterance. In that sense, the 'indeterminacy' of the potential that is extracted from utterances for future use is the whole point of having a human language rather than a pre-determined set of calls, cf. Deacon (1997): meanings in human languages are symbolic rather than situation-bound. Animal calls have determinate meanings because they can be reliably linked to definable situational circumstances – while human beings have the unique advantage of being able to draw on a system of meanings that can be flexibly re-applied in a variable range of situations.

This has consequences also for the question of what constitutes meaning, and what is "(non)semantic"; cf. the quotation above. It can be regarded as merely a terminological issue, but to reserve the key term 'meaning' for fully contextualized utterance meaning seems to me unfortunate for understanding the uniqueness of human language in enabling speakers to draw upon symbolic meaning. The goal of language description is sometimes said to be to describe everything a learner may need to know in order to use the language optimally. If the terminology above becomes accepted, however, this target state does not include knowing what words mean, or having knowledge of semantics – since 'meaning' and 'semantics' are bound up with concrete utterances that will never recur in exactly the same particular situations.

The input perspective can offer a useful supplement to the Croft and Cruse trajectory from vague purport to precise utterance meaning. In order to be a competent speaker, you have to know what constraints-on-variability the use of a word imposes on addressee understanding, and this knowledge can usefully be seen as constituting input-stage 'meaning'. Such constraints do not guarantee a particular reading, and thus they are (necessarily) indeterminate in relation to the output end; but viewed as input they can be surprisingly precise. In using e.g. the word 'computer' you invoke a fairly specific chunk of semantic substance and bring it to bear on whatever else you choose to say, and it seems to me that it would be practical to say that knowing it constitutes knowing the (input-stage) meaning of the word, even if there is some work to do before you get at the output-stage meaning.

Input 'meaning' emerges from the whole range of use, rather than being definable as invariant essence. Once it has emerged, however, it constitutes something else than merely

the sum total of actual usage events. As pointed out by Tyler and Evans (2003: 40), not all contextually discernible varieties of spatial configurations covered by *over* need to be specified as part of what is encoded by that word. Langacker's term 'centrality' (1987: 159) aptly captures what happens on the path of emergence from raw usage to (input) meaning: certain semantic properties are highlighted at the expense of others, and knowing a word entails knowing what features you centrally invoke when you use it. Centrality is thus a good way to allow for gradual and subtle differences. But it would seem desirable to also operate with more sharply defined differences: there are properties that are invoked and properties that are not invoked at all by given words. The meaning of the word *computer* is distinct from that of *grow* and that of *dirty* in ways that centrality alone cannot capture.

To sum up, there is a sense in which it would be absurd to say that input meanings are indeterminate: even in advance of a particular utterance, competent language users know what is (centrally) invoked by using a given word, precisely because, as ordinary people would say, they know the meaning of the word.

Another complementary perspective has to do with the doctrine of the precision of utterance meaning, which can also be given a different twist. Croft and Cruse (2004: 99) speak of a moment of 'crystallization' (analogous to recognizing a familiar face) that takes place when you understand a piece of language in use. While such experiences are central to understanding, it seems to me too strong a claim to suggest that they are an invariable concomitant of language understanding in context. The analogy appears to me most natural in the context of cases when you identify what is said with something that you have experienced on a precise previous occasion, and less obvious when you are struggling to understand something that is new to you (e.g. when someone is returning from a distant country and tries to make you understand what it was like). A comprehensive usage-based theory of meaning cannot take the achievement of such precision for granted. Output meanings, like input meanings, may be more or less vague. At both ends, most of the time we may expect something more than complete indeterminacy and settle for something less than blinding illumination.

Another problem for situating precise *word* meaning at the end stage is that even a precisely achieved situational understanding cannot be associated unambiguously with a single word. Croft and Cruse compare two occurrences of the expression 'raise one's head', one with a human subject and one involving a dog, and offer a compelling account of the differences in situated understanding. Both are said to be "a direct result of the occurrence of the word *raise*" (2004: 99). However, they are also the result of an act of composition, whereby the verb is combined with two different subjects – so the construals are bound up with the subject-verb-object combinations ('dog-raise-head' and 'woman-raise-head', roughly speaking) rather than specifically with the verb *raise*. More generally, utterance meanings are the result of the whole act, involving a combination of linguistic meanings as well as a combination of linguistic and non-linguistic aspects. To associate them directly with specific linguistic items, rather than via the speaker's 'extraction process', is therefore problematic.

For those reasons, a notion of meaning focused on the input stage appears to me more promising in specifying what human languages contribute to the whole process of language use in situated contexts. Processes moving from sprawling input meanings (such as those

associated with prepositions) to precise contextually determined senses are part of the picture. But so are language users' (including learners') processes of extracting word meanings from the sprawl of total usage events. Entrenched, conventional patterns and actual usage events mutually presuppose each other (cf. also Boye and Harder 2007).

5. 'Usage competence', compositionality, and the cooking analogy

As emphasized, I am not arguing that the input stage should be taken out of the context of actual language use. It is impossible to postulate semantic properties at the input end without having an answer to the question of what happens on the way to the end stage – which is the only one that we have direct intuitive access to. We know what we understand by an utterance, but we cannot follow the processes whereby words and context acting in concert get us to that stage. In order to be a cognitively and functionally competent language user, you have to possess the 'procedural' knowledge of how to get from the linguistic utterance to the relevant utterance understanding, drawing on all forms of context in the appropriate way.

There obviously needs to be an account of that stage as well. Ultimately a usage-based notion of processing competence must interface with all context types, including the 'language for special purposes' dimension; but in all cases the ability to take in linguistic input and use it to promote understanding is a shared feature. Attempts to be precise about linguistic-semantic properties viewed as input may be helpful in pursuing that aim. One way in which the instructional format may be conducive is that it specifies what the addressee is supposed to *do*, rather than operate with a purely representational format. The process leading from input to final understanding is bridged by the operations that have to be performed – as specified in the structural links encoded in a clausal utterance. Thus, if clausal structure is understood as a structured set of semantic instructions to be carried out by the addressee, the linguistic-semantic description interfaces directly with the processing stage.

An illustration of an instructional description of 'input' semantics can be given if we assume a version of the 'layered' scope hierarchy that is also inherent in Langacker's theory of the 'billiard ball model' and grounding predications, and that has been developed with great sophistication in Dik and Hengeveld-style Functional Grammar (Dik 1997; Hengeveld 2004) as well as in Role and Reference Grammar (Van Valin 1990). The main elements of clause meaning include verbs, arguments and grounding elements, organized in a hierarchy where the 'billiard ball model', i.e. the structure consisting of the verb and participants, is at the bottom, inside the scope of tense and, at the top, illocutionary type. A rough sketch of the semantic structure of the simple clause *the test happened* would be

declarative (past tense (HAPPEN (definite (TEST)))

Rephrased in the instructional format, the coding 'programme' that an addressee faces could be rendered (bottom-up) as

Identify the relevant test (cf. above) and predicate the event type 'happen' about it; understand the resulting predication as applying to the relevant past situation (past tense), and take it as a fact (declarative).

This format, although it involves bracketed structure, is not specific to linguistic competencies. It can be used also for cooking recipes (cf. Harder 1996: 214):

Serve (sprinkle with lemon (grill (add salt and pepper (slice (salmon))))))

The bottom-up paraphrase is "take a salmon, slice it, add salt and pepper, put the slices on the grill and sprinkle them with lemon before serving". There is one more analogy between linguistic utterances and cooking recipes beyond their status as structured instructions: at the bottom of the hierarchy we find the more 'substantial' elements, while the higher elements are more obviously 'procedural'. The general format is that you achieve results by taking the 'raw materials' as it were, subsequently performing procedures on them. In the case of language, the substantial raw materials are the 'billiard-ball elements', the arguments and predicates that are accessed via lexical meanings-cum-cognitive models (in the mind) or deictic meanings (from the situational context) – and these are then used and combined in ways specified by the higher-level elements.

In order to be able to handle such structured sets of instructions, addressees must possess the necessary competencies; and one aspect of those is the input dimension: they must be able to grasp what is required of them. An interesting property of the instructional stage is that the structure is fully compositional: the operations must occur in the specified order. That does not guarantee all properties of the final output – but unless the input is understood and performed in an adequately determinate way, an acceptable output would be pure fluke.

Croft and Cruse (2004: 105) also invoke the cooking analogy in discussing the issue of compositionality and ask, rhetorically, "is cookery a compositional art?" Because they apply the analogy to the output stage, their answer reflects the 'partial compositionality' that is generally recognized in cognitive linguistics: although 'red hats' may be the intersection of the sets of 'red things' and 'hats', not all semantic properties can be expected to be as well behaved. But if we view compositionality in relation to the instructional stage, full compositionality is achievable to the extent that the input expressions are independent entities. Even recalcitrant cases like 'fake' work compositionally at the input stage: *fake Japanese paintings* are different from *Japanese fake paintings* precisely because the operation of turning them into fakes occurs at different points in the encoded instructional procedure (cf. Harder 2003: 94). A human language user must be up to the compositional logic in order to understand what is going on.

Similarly for cooks: they must be able to factor out the operations they need to perform and compose them correctly. Unless they are capable of frying at the right point and for the right time, or adding the right quantity of salt at the right time, the final product will suffer – even if at that final 'output' stage the results of frying and adding salt are everywhere and can no longer be factored out from the total product.

6. Final remarks

I have argued that a comprehensive usage-based account of language needs to be precise about the semantic contribution of the linguistic code to actual utterance situations. The reason of principle is that as linguists we are committed to being as precise as we can about our object of description. A more practical reason is that some properties can only be captured if we concentrate on the stage of the input to the utterance situation, rather than the final output. In aiming to be precise about the input, we simultaneously further a process of exploring usage-based competence, i.e. what it is a human language user needs to be able to do with a linguistic input (the means) in order to get at the utterance understanding (the end). Reasons for underplaying such a distinction between the linguistic input and the total process may include fear of backsliding to a view of language as an autonomous object – but in fact such a distinction is a prerequisite for understanding the way language is inherently adapted to link up with cognitive and situational factors. In bringing this point to bear on current trends in exploring the usage-based perspective, I have argued that reserving the term ‘meaning’ so that it applies to actual utterance meanings only would be detrimental to the project of understanding what exactly language can do for you.

References

- Allwood, J. 2003. Meaning potentials in context. In *Cognitive approaches to lexical Semantics*, H. Cuyckens, R. Dirven & J. R. Taylor (eds), 29–65. Berlin: Mouton de Gruyter.
- Boye, K. & Harder, P. 2007. Complement-taking predicates: usage and linguistic structure. *Studies in Language*.
- Clark, H. H. 1998. Communal lexicons. In *Context in Language Learning and Language Understanding*, K. Malmkjær & J. Williams (eds), 63–87. Cambridge: Cambridge University Press.
- Contini-Morava, E., Kirsner, R. S. & Rodriguez-Bachiller, B. (eds). 2004. *Cognitive and Communicative Approaches to Linguistic Analysis*. Amsterdam: John Benjamins.
- Croft, W. A. 2000. *Explaining Language Change. An Evolutionary Approach*. London: Longman.
- Croft, W. A. 2001. *Radical Construction Grammar. Syntactic Theory in Typological Perspective* Oxford: University Press.
- Croft, W. A. & Cruse, D. A. 2004. *Cognitive Linguistics*. Cambridge: Cambridge University Press.
- Croft, W. A. & Poole, K. T. 2008. Inferring universals from grammatical variation: multidimensional scaling for typological analysis. *Theoretical Linguistics* 34: 1–37.
- Davies, J. M. D. & Isard, S. D. 1972. Utterances as programs. In *Machine Intelligence*, D. Michie & B. Meltzer (eds), 325–339. Edinburgh: Edinburgh University Press.
- Deacon, T. 1997. *The Symbolic Species. The Co-Evolution of Language and the Human Brain*. Harmondsworth: Penguin.
- Dik, S. 1997. *The Theory of Functional Grammar*. Berlin: Mouton de Gruyter.
- Ducrot, O. 1972. *Dire et ne pas dire. Principes de sémantique linguistique*. Paris: Hermann.
- Engberg-Pedersen E., Fortescue, M., Harder, P., Jakobsen, L. F. & Heltoft, L. (eds). 1996. *Content, Expression and Structure. Studies in Danish Functional Grammar*. Amsterdam: John Benjamins.
- Evans, V. 2006. Lexical concepts, cognitive models and meaning-construction. *Cognitive Linguistics* 17 (4): 491–534.
- Fauconnier, G. [1985] 1994. *Mental Spaces*. Cambridge: Cambridge University Press.

- Harder, P. & Kock, C. 1976. The Theory of Presupposition Failure. (Travaux du cercle linguistique de Copenhague XVII.) København: Akademisk forlag.
- Harder, P. & Togeby, O. 1993. Pragmatics, cognitive science and connectionism. *Journal of Pragmatics* 20: 467–492.
- Harder, P. 1996. *Functional Semantics. A Theory of Meaning, Structure and Tense in English*. Berlin: Mouton de Gruyter.
- Harder, P. 2003. Mental spaces: exactly when do we need them? *Cognitive Linguistics* 14 (1): 91–96.
- Heim, I. 1983. File change semantics and the familiarity theory of definiteness. In *Meaning, Use, and Interpretation of Language*, R. Bäuerle, C. Schwarze & A. von Stechow (eds), 164–189. Berlin: Mouton de Gruyter.
- Hengeveld, K. 2004. The architecture of a functional discourse grammar. In *A New Architecture for Functional Grammar*, J. Mackenzie, J. Lachlan & M. A. Gómez-González (eds), 1–21. Berlin: Mouton de Gruyter.
- Hjelmslev, L. 1943. *Omkring sprogteoriens grundlæggelse*. København: Københavns Universitet. (English Translation, 1953. *Prolegomena to a Theory of Language*. Bloomington: Indiana University)
- Lakoff, G. & Johnson, M. 1999. *Philosophy in the Flesh. The Embodied Mind and its Challenge to Western Thought*. New York: Basic Books.
- Langacker, R. W. 1987. *Foundations of Cognitive Grammar, Volume 1: Theoretical Prerequisites*. Stanford: Stanford University Press.
- Langacker, R. W. 1991. *Foundations of Cognitive Grammar, Volume 2: Descriptive Applications*. Stanford: Stanford University Press.
- Nølke, H., Fløttum, K. & Norén, C. 2004. *ScaPoLine – La théorie scandinave de la polyphonie linguistique*. Paris: Éditions Kimé.
- Russell, B. [1905] 1956. On denoting. In *Logic and Knowledge*, R. C. Marsh. (ed.), 39–56. London: Macmillan.
- Searle, J. R. 1969. *Speech Acts*. Cambridge: Cambridge University Press.
- Sperber, D. & Wilson, D. 1986. *Relevance: Communication and Cognition*. Oxford: Blackwell.
- Swinney, D. A. 1979. Lexical access during sentence comprehension: (re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior* 18: 523–534.
- Thompson, S. A. 2002. ‘Object complements’ and conversation. *Studies in Language* 26 (1): 125–164
- Tyler, A. & Evans, V. 2003. *The Semantics of English Prepositions. Spatial Scenes, Embodied Meaning and Cognition*. Cambridge: Cambridge University Press.
- Van Valin, R. D. 1990. Layered syntax in role and reference grammar. In *Layers and Levels of Representation in Language Theory. A Functional Approach*, J. Nuyts, A. M. Bolkestein & C. Vet (eds), 193–231. Amsterdam: John Benjamins.

Semantic representation in LCCM Theory

Vyvyan Evans

1. Introduction

In this paper I am concerned with the approach to semantic representation adopted in the *Theory of Lexical Concepts and Cognitive Models*, or *LCCM Theory* for short.¹ LCCM Theory takes its name from the two central theoretical constructs adopted in the theory: the *lexical concept* and the *cognitive model*. A lexical concept is a component of linguistic knowledge, the semantic pole of a *symbolic unit* (in Langacker's e.g. 1987 terms), and encodes a bundle of different types of linguistic knowledge. Put another way, the lexical concept represents the means adopted in LCCM Theory of modelling units of *semantic structure*. In contrast, a cognitive model is a component of conceptual knowledge, which is to say, non-linguistic knowledge. Hence, the cognitive model represents the means adopted in LCCM Theory of modelling units of *conceptual structure*. LCCM Theory assumes that lexical concepts and cognitive models are types of semantic representation belonging to two distinct representational systems, which have distinct and divergent functions. These are the *linguistic system*, which encodes semantic structure and the *conceptual system* which encodes conceptual structure. Following arguments presented by Barsalou et al. (to appear), I suggest that the linguistic system evolved, in part, by facilitating more effective control of the extant representations in the conceptual system. That is, linguistic representations are specialised for providing a 'scaffolding' to structure conceptual representations, thereby facilitating their use in communication (cf. Talmy 2000). While the much older – in evolutionary terms – conceptual system evolved for action and perception, i.e. for non-linguistic purposes, the emergence of language facilitated the use of conceptual representations in linguistically-mediated meaning construction, thereby providing modern humans with a significant evolutionary advantage. With the association of linguistic and conceptual representations, humans were able to engage in the advanced symbolic behaviours that led to the explosion of sophisticated ritual practice, material culture, art and science around 50,000 years ago during the period known as the Upper Palaeolithic (Mithen 1996).

My argument, in a nutshell, is this: the semantic representations in the linguistic and conceptual systems interact for purposes of linguistically-mediated communication. Together, the lexical concept and the cognitive model form a level of representation that I

1. LCCM Theory represents an attempt to provide a cognitively realistic account of the semantic mechanisms involved in the construction of meaning in language understanding.

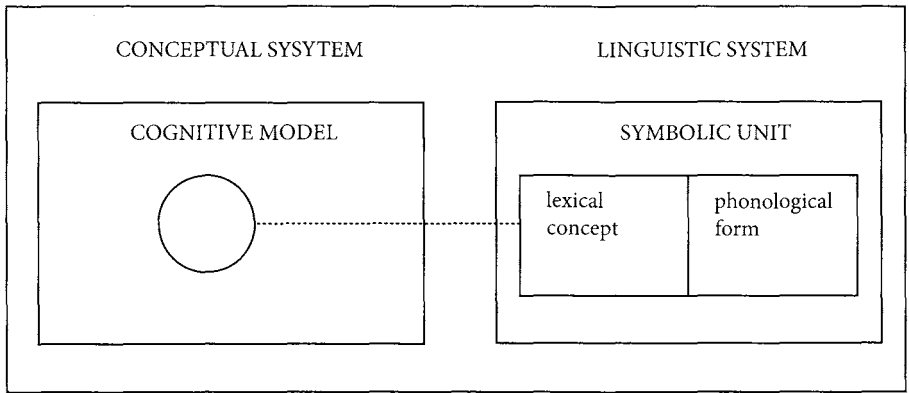


Figure 1. The relationship between the linguistic and conceptual systems

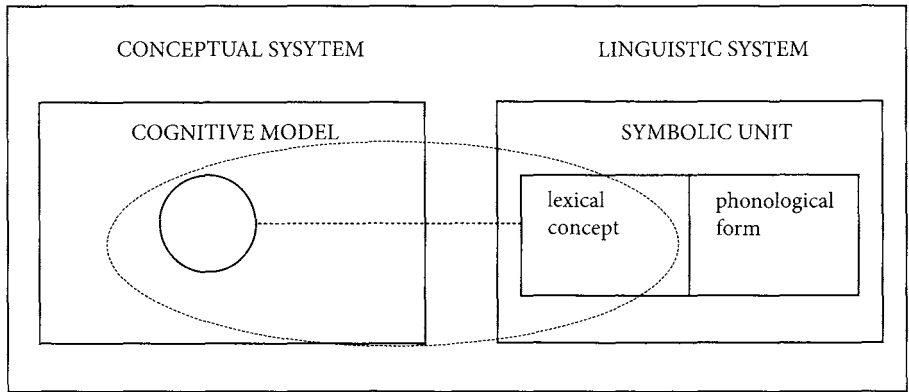


Figure 2. Semantic representation in LCCM Theory

refer to as *semantic representation*. My purpose in this paper is to describe the nature of the lexical concept, the nature of the cognitive model, and the nature of the interaction between the two.

In order to better illustrate these distinctions, and how they intersect, Figure 1 provides a diagrammatic representation of the relationship between the linguistic and conceptual systems, as assumed by LCCM Theory. Figure 2 provides an illustration of semantic representation in LCCM Theory. In Figure 1 the dashed line between the lexical concept in the linguistic system and the cognitive model (represented by the circle) in the conceptual system represents a path of access relating the two. Figure 2 is the same as Figure 1 except that it additionally features a dashed ellipse encircling the lexical concept (in the linguistic system) and the cognitive model (in the conceptual system), the two types of representations which collectively comprise semantic representation.

2. The distinction between linguistic content and conceptual content

According to Talmy (2000), a central design feature of language is that the concepts expressed are divided into two subsystems. Talmy characterises this in terms of what he refers to as the *grammatical subsystem* and *lexical subsystem*. These two subsystems serve to express what we might refer to as the experiential complex – Talmy uses the term: *cognitive representation* – that a speaker attempts to evoke in the listener by virtue of deploying language. The range of concepts expressed by the grammatical subsystem is highly restricted cross-linguistically, providing a basic framework for the structuring of the experiential complex that language users seek to evoke in their interlocutors. Put another way, the lexical concepts associated with the grammatical subsystem have *schematic content*, providing a structuring function. Thus, lexical concepts of this sort provide a ‘scaffolding’ so to speak, across which the rich content associated with lexical concepts of the lexical subsystem can be draped. In contradistinction to this, the lexical concepts associated with the so-called lexical subsystem provide *rich content*, giving rise to the details (rather than structural aspects) of the cognitive representation.

An important aspect of Talmy’s work is the claim that the distinction between rich versus schematic content corresponds to a bifurcation between form types: open-class versus closed-class forms. Closed-class forms are so-called because it is considered more difficult to add members to this set. This set of lexical items includes the so-called ‘grammatical’ or ‘function’ words such as conjunctions, determiners, pronouns, prepositions, and so on. In contrast, open-class forms include words belonging to the lexical classes: noun, verb, adjective and adverb.

While the concepts expressed by closed-class forms encode schematic content, they are nevertheless essential for the expression of the cognitive representation. To make this point clear, consider the following semantic analysis of the range of open- and closed-class elements which comprise the utterance in (1):

- (1) *A waiter served the customers*

The forms in bold: **a**, **-ed**, **the** and **-s** are associated with the grammatical subsystem. Their semantic contribution relates to whether the participants (waiter/customers) in the experiential complex evoked by (1) can be easily identified by the hearer (the use of the indefinite article **a** versus the definite article **the**), that the event took place before now (the use of the past tense marker **-ed**), and how many participants were involved (the absence or presence of the plural marker **-s**).

In contrast, the forms in italics: *waiter*, *serve* and *customer* are associated with the lexical subsystem. That is, their semantic contribution relates to the nature of participants involved in the experiential complex, and the relationship holding between them. In other words, while the closed-class forms encode content relating to structural aspects of the experiential complex evoked, the open-class forms are associated with detailed

information concerning the nature of the participants, scenes involving the participants, and the states and relationships that hold.²

I argue that the distinction in content evoked by language, and pointed to by Talmy, relates to a distinction in content associated with the linguistic system (and lexical concepts) on the one hand and the conceptual system (and cognitive models) on the other. The two distinct types of content implicated I refer to as *linguistic content* and *conceptual content*.

Dealing with the latter first, conceptual content relates to the rich content evoked by open-class forms. Information of this kind is primarily perceptual in nature. By 'perceptual' I have in mind information that derives from i) *sensory-motor systems*—those sensory systems that recruit information relating to the external environment and the human individuals' interaction with the environment – as well as ii) *proprioception* – the systems that recruit information relating to the motor aspects of the body's own functioning – and iii) *subjective experience* – which includes experiences ranging from emotions, temporal and other cognitive states, to the visceral sense, all discussed in more detail later. Accordingly, I am following Barsalou (1999) in defining perceptual experience more broadly than has traditionally been the case.

Conceptual content provides records of perceptual states, in the sense just given. As such, it is *analogue* in character: it re-presents perceptual information that parallels the multimodal perceptual experience that it constitutes a representation of.³ As such, conceptual structure is not suitable for being encoded in language. After all, language as a representational system, consisting of symbolic units, is simply not equipped to directly encode the rich, multimodal aspects of perceptual experience.

In contrast, I argue that the schematic content identified by Talmy is not an analogue representation of perceptual experience. Rather, it represents an abstraction over perceptual content of various sorts, provided in a form that can be encoded in language, i.e. by lexical concepts. Content of this kind constitutes what I refer to as linguistic content, and forms part of the information encoded by a lexical concept.

While the distinction between rich and schematic aspects of the cognitive representation provides the basis for my distinction between linguistic and conceptual content, the distinction in open-class and closed-class forms provides evidence for a closely related distinction concerning the nature of lexical concepts – recall that lexical concepts are conventionally associated with phonological vehicles (i.e. forms). The distinction in vehicle types provides evidence that lexical concepts fall into two distinct categories. Closed-class vehicles are associated with lexical concepts which are specialised for encoding linguistic content. Lexical concepts of this sort I refer to as *closed-class lexical concepts*. Open-class

2. The closed-class forms mentioned thus far all have an overt phonetic realisation. However, each of the examples discussed also include closed-class forms that are phonetically implicit. Examples of phonetically implicit forms include lexical classes: e.g. noun, verb; lexical subclasses: e.g. count noun, mass noun; grammatical relations: e.g. subject, object; declarative versus integrative forms, active voice versus passive voice, and clause-level symbolic units such as the ditransitive construction, and so forth.

3. Conceptual content is not an exact record of the perceptual states that are captured. Rather, it is somewhat attenuated. See Barsalou (1999) for discussion.




Figure 3. The distinction in content associated with lexical concepts

vehicles, while also encoding linguistic content, are, in addition, specialised for serving as *access sites* to conceptual content. Lexical concepts of this sort I refer to as *open-class lexical concepts*. This distinction is captured in Figure 3.

The distinction between the terms ‘encode’ and ‘afford access’ in the previous paragraph is critical here. Linguistic content is encoded by lexical concepts precisely because this is the content which makes up lexical concepts. However, conceptual content is associated with a different representational type, the cognitive model, which is non-linguistic in nature. Thus, conceptual content is not directly encoded in language, although the linguistic system has developed the means to access conceptual content, discussed in more detail later in the paper. Table 1 provides a summary of the way some of the key terms introduced so far are used in LCCM Theory.

Table 1. A summary of key terms in LCCM Theory

Term	Description
Linguistic system	The collection of symbolic units comprising a language, and the various relationships holding between them
Symbolic unit	A conventional pairing of a phonological form or vehicle and a semantic element
Lexical concept	The semantic element that is paired with a phonological vehicle in a symbolic unit
Linguistic content	The type of content encoded by a lexical concept. This content is of a highly schematic type that can be directly encoded in language
Conceptual system	The body of non-linguistic knowledge captured from perceptual experience that is made of perceptual states. This knowledge derives from sensory-motor experience, proprioception and subjective experience
Cognitive model	The representational form that knowledge in the conceptual system takes, as modelled in LCCM Theory. Consists of frames which give rise to a potentially unlimited set of simulations
Conceptual content	The nature of the knowledge encoded by a cognitive model
Lexical representation	The primary substrate deployed in linguistically-mediated meaning construction, and modelled in terms of symbolic units and cognitive models
Semantic representation	The semantic dimension of lexical representations, consisting of semantic structure and conceptual structure
Semantic structure	That part of semantic representation encoded by the linguistic system. Semantic structure is modelled, in LCCM, Theory, by lexical concepts,
Conceptual structure	That part of the semantic representation encoded by the conceptual system. Conceptual structure is modelled, in LCCM Theory, by cognitive models

3. An illustration

Before proceeding further, I provide a brief illustration of the distinction between linguistic and conceptual content. Consider the use of the lexical item *red* in the following examples, adapted from Zwaan (2004):

- (2) a. The teacher scrawled in red ink all over the assignment
 b. The red squirrel is in danger of becoming extinct in the British isles

Zwaan makes the point that in linguistic examples such as (2), *red* designates two different sorts of sensory experience. That is, while the hue derived from the use of *red* in (2a) is quite a vivid red, the hue prompted for by (2b) is likely to be closer to a dun/browny colour. That is, what I refer to as the *semantic potential* of *red* is not 'there' in the word itself. That is, whatever *red* designates, we are not dealing with purely linguistic knowledge. Rather, the word *red* provides access to perceptual information and knowledge, which can be reconstructed or *simulated* – I will have more to say about this idea below. Put another way, the hue derived is not a function of linguistic knowledge, but relates to what I am referring to as conceptual content. This is not to say that *red* does not provide linguistic knowledge. In point of fact the vehicle *red* has an associated lexical concept that I gloss as [RED] – in LCCM Theory lexical concepts are glossed by a term, here 'red', in small capitals inserted in square brackets. This encodes schematic linguistic content, signalling that an entity is being referred to, that the entity being referred to is a relation of some kind, and that the relation is specifically an attribute of a thing. In short, while linguistic content includes highly schematic semantic knowledge, conceptual content concerns richly detailed perceptual knowledge.

4. The nature of semantic structure

In LCCM Theory semantic representation consists of units of semantic structure and conceptual structure and their interaction. In this section I address, in more detail, the nature of semantic structure.

Semantic structure is modelled in terms of the theoretical construct of the lexical concept, which constitutes a unit of semantic structure. As noted above, lexical concepts encode linguistic content. Linguistic content represents the informational form that conceptual structure takes for *direct* representation *in* language. Put another way, linguistic content takes a form that can be encoded in a format that is externalised in an auditory stream (or a manual gestural stream in the case of signed language), which is time-pressured. Such a format presumably requires filtering out the complexity associated with the range of perceptual experiences – in the sense defined above – encoded. Hence, a lexical concept can be thought of as a bundle of different types of highly schematic content which is thereby specialised for being encoded in language. I detail some of the key aspects of linguistic content below.

4.1 Parameters

One way in which knowledge, in general terms, can be represented is in terms of richly inflected nuances that serve to reflect the complexity of experience. An alternative way is to 'compress' such fine distinctions into two, three or more, much broader, and hence, far more general distinctions. These I refer to as *parameters*. Linguistic content serves to encode content by adopting the latter strategy, which is to say, to employ *parameterisation*. Parameters are hence part of the bundle of information that a lexical concept serves to encode.

To illustrate this notion, consider the complex range of expressions that a language user might employ, in English, in order to 'locate' themselves with respect to time, thereby facilitating time-reference. Any one of the following could conceivably be employed, depending upon context: *today, January, 2008, the day after yesterday, the day before tomorrow, this moment, now, this second, this minute, this hour, today, this week, this month, this quarter, this year, this half century, this century, this period, the 8th day of the month, this era, this millennium*, and so on.

In contrast, parameterisation functions by dividing all the possible permutations relating to a given category, such as time-reference, into a small set of divisions: parameters. Such parameters might distinguish between the past, for instance, and the non-past. Indeed, this is the basis for the tense system in English, as illustrated by the following:

- | | | | |
|-----|----|--------------------|----------|
| (3) | a. | He kicked the ball | Past |
| | b. | He kicks the ball | Non-past |

English encodes just two parameters that relate to Time-reference: Past versus Non-past, as exhibited by the examples in (3), and thus manifests a binary distinction. Some languages, such as French, have three parameters: Past, Present and Future. Some languages have more than three parameters, distinguishing additionally remote past from recent past, for instance. The language with the most parameters for time-reference thus far reported is an African language: Bamileke-Dschang with eleven. Crucially, parameters are encoded by specific lexical concepts, and thus form part of the knowledge bundle that constitutes a lexical concept. For instance, the parameter Past is encoded, in English, by the lexical concept associated with the *-ed* form in (3a). However, other lexical concepts also include the parameter Past such as the lexical concepts associated with the following forms: *sang, lost, went*, etc.

I argue, then, that a key feature of linguistic (as opposed) to conceptual content is that it encodes knowledge in parametric fashion. Parameterisation is a highly reductive form of abstraction: it serves to abstract across the complexity exhibited by a particular category. In consequence the parameters encoded by linguistic content serves to 'strip away' most of the differences apparent in the original experience, thereby reducing it to a highly limited numbers of parameters.

4.2 The non-analogue nature of linguistic content

As conceptual content relates to records of perceptual states captured directly from a variety of experience types including sense perception, proprioception and subjective

experience, it therefore consists, as noted above, of perceptual states recorded in analogue fashion: in a format that is similar to the perceptual experiences that gave rise to them. Indeed, there is a good deal of evidence, in the neuroscience literature, that sensory-motor representations, for example, are stored in the same areas of sensory-motor cortex that process sensory motor experience (Pulvermüller 1999).

In contrast, I argue that linguistic content is so highly schematic in nature that it is *non-analogue*: it takes a format that is not analogous to the perceptual experiences that it is a schematisation of. Hence, due to the reduction of rich perceptual information to highly impoverished parameters, this gives rise to a qualitatively very different type of information from the kind captured by conceptual content. To illustrate, re-consider the parameters Past and Non-past discussed with respect to example (3) above. These parameters are highly schematic abstractions drawn from the complex range of temporal relationships that hold between our experience of past, and our experience of now: our temporal location as experiencing centres of consciousness. Temporal experience, a form of subjective experience, is extremely rich in perceptual terms (Evans 2004a). Yet the parameters Past and Non-past are not rich at all.

In sum, parameters encode highly schematic linguistic content abstracted from far richer perceptual experience, as recorded in the conceptual system, and provide a means for encoding recurrent ‘digitised’ dimensions of humanly relevant experience in an efficient way. In contrast, conceptual content which is accessed via open-class lexical concepts, gives rise to perceptually rich aspects of experience, about which I will have more to say later.

4.3 Topological reference

A further consequence of the highly reductive nature of the parameters encoded as linguistic content, and one first pointed to by Talmy (e.g. 2000), is that they provide *topological reference* rather than *Euclidean reference*. That is, linguistic content encodes schematic aspects of sensory-motor, proprioceptive and subjective experience, while conceptual content, to which open-class lexical concepts facilitates access, relates to precise, metric distinctions.

To illustrate consider the closed-class lexical concepts associated with the demonstrative vehicles *this* and *that*. The lexical concepts associated with these vehicles encode a distinction between an entity construed as proximal to the speaker, glossed as [THIS], versus an entity construed as distal, glossed as [THAT]. The distinction between the lexical concepts [THIS] versus [THAT] is illustrated by (4):

- (4) “Sit on this chair not that one!”

In this utterance, the chair that the addressee is being asked to sit on is the one closer to the speaker: ‘this chair’ as opposed to ‘that one’. Nevertheless, the distinction between [THIS] versus [THAT] does not rely upon precise metric details such as the exact distance from the speaker, in terms of metres, centimetres and millimetres, for instance. After all, it is immaterial how far the chairs are from the speaker (within reason), as long as one is closer to the speaker than the other. In other words, closed-class lexical concepts are *magnitude*

neutral, where magnitude has to do with metric properties relating to distance. This is what it means to say that closed-class lexical concepts provide topological reference.

In contrast, the open-class lexical concepts, in addition to encoding linguistic content also, additionally, facilitate access to conceptual content, and hence can be employed to express metric details of distance giving rise to Euclidean reference. This is illustrated by (5):

- (5) “Sit on the chair 2.54 metres away from me!”

The expression ‘2.54 metres’ involves open-class lexical concepts rather than closed-class lexical concepts, and serves to evoke with greater precision the chair in question.⁴

4.4 A restricted set of domains and categories

A consequence of parameterisation is that the range of *domains*, and the member *categories* that populate them, are highly restricted in terms of their encoding as parameters in linguistic content (cf. Talmy 2000). In using the term domain I have in mind large-scale and coherent bodies of knowledge such as the following: TIME, SPACE, COLOUR, MOTION, FORCE, TEMPERATURE, MENTAL STATES, and so on. By category I have in mind the member notions that populate a particular domain. For instance, in terms of the domain of TIME, categories consist of notions such as Punctuality, Durativity, Sequentiality, Simultaneity, Synchronicity, Boundedness, Time reference (e.g. Past, versus Non-past etc.), Time-reckoning (e.g. 10.05pm, etc.), and so forth. While *all* the domains of the sort just mentioned, and the categories which populate them, are evident at the conceptual level, only a restricted subset are encoded at the linguistic level, in terms of linguistic content.

For instance, some domains to which open-class lexical concepts facilitate access, such as COLOUR, do not appear at all in terms of linguistic content in English or any other language. That is, there are no parameters, in the sense defined above, that relate to this domain. This follows as many (perhaps most) domains do not relate to experience that can be straightforwardly parameterised in a humanly relevant way. There are at least two likely explanations for this. Firstly, the nature of the domain in question may not lend itself to being ‘reduced’ to highly schematised digitised parameters. After all, the reduction to content that does not directly give rise to simulations results in a reduction that, for some domains such as COLOUR, may eliminate the essential character of the information thereby making it uninterpretable. A second reason is that some domains do not relate in a ubiquitous way to the humanly relevant scenes that language serves to encode. For instance, categories that relate to the domain of MEDIAEVAL MUSICOLOGY, or even categories that relate to less esoteric domains such as LOVE or JOURNEYS are not as ubiquitous in human experience as categories relating to domains such as SPACE, TIME, MOTION and MENTAL STATES which do appear to be encoded in linguistic content.

4. The parameters encoded as linguistic content exhibit a range of other *Euclidean neutralities*: notably with respect to the domains of space and time. See Evans (to appear a) for details.

In addition to the restricted set of domains encoded, linguistic content also features only a small number of categories within each domain. To illustrate, consider a few of the categories associated with the domain *TIME*:

Domain: <i>TIME</i>		
	Category: Time-reference	Parameter
(6)	a. He kicked the ball	Past
	b. He kicks the ball	Non-past
	Category: Boundedness	
(7)	a. Holly has left the party	Bounded
	b. Holly is leaving the party	Unbounded
	Category: Plexity	
(8)	a. Fred coughed	Uniplex
	b. Fred coughed for 10 minutes	Multiplex

The category that I refer to as Time-reference is more traditionally referred to as tense. Of the other two categories illustrated, these are normally treated as relating to what is commonly referred to as aspect. The examples in (7) are usually referred to as perfective and imperfective aspect. The more usual terms for Uniplex and Multiplex in (8), as they relate to *TIME*, are ‘semelfactive’ and ‘iterative’ respectively.

4.5 Nominal versus relational lexical concepts

Another aspect of linguistic content is that it encodes a bifurcation between *nominals* and *relations* (Langacker 1987). The distinction in type of lexical concepts is as follows. Nominal lexical concepts are *conceptually autonomous*: they relate to entities which are independently identifiable, such as ‘chair’, or ‘shoe’. In contrast, relations are *conceptually dependent*: they constitute a relation holding between other entities, and are thus ‘dependent’ on those other entities in order to fully determine the nature of the relationship. For instance, in an utterance such as the following:

- (9) Max hid the mobile telephone under the bed.

The lexical concept associated with the form *hid*, which I shall gloss as [HID], relates the conceptually autonomous lexical concepts associated with the vehicles *Max*, *mobile telephone* and *bed*, establishing a relationship involving ‘hiding’ between the conceptually autonomous participants in the linguistically-mediated *conception*:⁵ namely [MAX] and [BED]. Analogously, the lexical concept associated with the vehicles *under* establishes a spatial relation between lexical concepts associated with *mobile telephone* and *bed*.

The conceptually dependent structure of relational lexical concepts is modelled, in LCCM Theory, in terms of a schematic *participant role* (Goldberg 1995). The lexical con-

5. ‘Conception’ is a technical term used in LCCM Theory to refer to utterance meaning.

cept [HID] as exemplified in (9) encodes three schematic participant roles.⁶ The rich content relating to the participant roles is not specified in linguistic content. This arises from access to conceptual structure. That is, conceptual structure encodes rich content relating to hiding: that it involves someone that does the hiding for particular reasons, and that an entity of a particular sort, often an object, is hidden. Non-linguistic knowledge also includes what facilitates something being hidden, such as perceptual inaccessibility of the object being hidden and/or its being placed in a novel location. Conceptual content also includes information relating to the motor processes involved in hiding, which involve moving the object from one location to another. The participant roles encoded as part of the linguistic content for [HID] do not encode such details. Rather, what is encoded is a highly abstract representation, derived from the rich perceptual details of a hiding scenario. As such we have three roles that serve to distinguish between the three entities involved at the most general level of detail. These participant roles are: Hider, Object and Location.

Just as the bifurcation in lexical concepts discussed above – that holding between lexical concepts which solely encode linguistic content and those which additionally facilitate access to conceptual content – corresponds to a distinction in the formal encoding of lexical concepts: the distinction between open and closed-class vehicles, so too the distinction between nominal and relational lexical concepts has a formal reflex in terms of linguistic vehicles. In a language such as English, for instance, this distinction relates to lexical concepts associated with what are commonly referred to as nouns and noun phrases (nominals) on the one hand, and lexical concepts associated with other lexical forms, including verbs, prepositions, adjectives, adverbs and non-finite verb forms such as infinitives and participles (relations) on the other (see Langacker 1987 for details).

4.6 Referentiality

Another key aspect of linguistic content is that it is inherently referential in nature. Referentiality takes a number of different forms, as detailed below. However, the defining feature is that lexical concepts serve to encode the following: an intention that a particular entity is being indexed or, more informally, 'pointed to'. In using the term 'entity' I have in mind physical entities that inhabit the world such as people, as well as physical artefacts, such as 'Sam' and 'ball' in (10a), abstract notions such as ideas, for example 'peace' in (10b), as well as relations that hold between physical entities and abstract ideas, such as 'kicked' in (10a) and 'thought about' in (10b), as well as highly schematic relations, as encoded by 'to' in (10c).

- (10) a. Sam kicked the ball
 b. Sam thought about peace
 c. Sam walked to the park

I identify at least three distinct types of *reference* encoded by lexical concepts.

6. Notice that the form *hid* is polysemous. For instance, *hid* is also associated with the 'reflexive' lexical concept in which an entity hides oneself, as in: *John hid in the wardrobe*. This lexical concept, which I gloss as [REFLEXIVE HID] encodes two schematic participant roles.

The first type relates to what I will refer to as *denotational reference*. Many lexical concepts serve to index a physical entity of some sort, whether real or imagined. In this sense, part of what the lexical concepts associated with the vehicles *John* and *unicorn* serve to do is to signal an intention, on the part of the speaker, to refer to a given entity.

The second type I refer to as *cognitive reference*. This relates to relatively abstract notions or ideas that have no physical substance, whether real or imagined, and relate to lexical concepts associated with forms such as *love*, *war*, *phonology*, and so forth. Hence, lexical concepts that serve to encode cognitive reference serve to signal an intention, on the part of the speaker, to refer to a non-physical idea.

The third type I refer to as *contextual reference*. This involves reference to an entity that is present in the linguistic or extra-linguistic discourse context. Hence, reference of this sort involves the encoding, by a lexical concept, of an intention to refer to an entity that the addressee can recover from context.

One type of contextual reference is *textual reference*. One form of textual reference involves reference to an entity already mentioned. This is traditionally termed *anaphora*. Textual reference that relates to an entity yet to be mentioned is termed *cataphora*. Examples of textual reference are provided in the examples below.

- (11) a. John is smart. He had a reading age of 14 by the time he was just 8.
 b. I want to say just this: I love you.
 c. The new target to reduce carbon emissions by 20% by 2020 will be a tough thing to achieve.

In the examples in (11a), the lexical concepts associated with the forms *he*, *this* and *thing* are specialised for referring to other entities (underlined) in the text.

There are many kinds of lexical concepts which encode an intention to signal contextual reference as it relates to extra-linguistic context. Many of these are often treated under the heading of *deixis*. Previous research has identified a range of diverse sorts of deictic lexical concepts including phenomena referred to as *spatial deixis*, *temporal deixis* and *social deixis* (for details see Fillmore 1997; Levinson 1983).

4.7 Pragmatic point

The final dimension of linguistic content that I address here relates to what I refer to as pragmatic point. This is a term I borrow from Fillmore et al. (1988). I use this term to refer to schematic aspects of the extra-linguistic dimensions of the encoding of linguistic content by a given lexical concept. As I use it, this term relates, broadly, to two aspects: i) the contexts of use in which a given lexical concept is conventionally employed, including settings and participants, and ii) some aspects of what has traditionally been referred to as the *illocutionary point* (Searle 1969) of a given lexical concept: which is to say the communicative purpose for which a lexical concept is employed.⁷

7. It is worth re-emphasising here that linguistic content is schematic in nature. Hence, while making a *speech act* (Searle 1969), such as declaring a state of war, for example, involves being able to call upon highly detailed bodies of conceptual knowledge relating to the sorts of scenarios and participants involved,

To illustrate the notion of pragmatic point consider the vehicle *declared* in the examples below. This is associated with at least three English lexical concepts, each of which exhibits a different pragmatic point.

- (12) a. She declared her love for him
 b. Chamberlain declared war on Germany on September 3rd 1939
 c. Despite being over the limit on the amount of dollars in cash eligible to be taken into the country, she declared nothing as she crossed the US border.

The use of *declared* in (12a) serves to encode an intention to provide information of a particular sort, with an above-average level of assertiveness. Hence, the lexical concept which sanctions this use of *declared* can be glossed as [FORTHRIGHT INFORMATIONAL ASSERTION]. In contrast, the lexical concept associated with the use of *declared* in (12b) relates to an assertion which either changes, or otherwise revises, an institutional state. Crucially, not only is the illocutionary point distinct from the lexical concept responsible for the use of *declared* in (12a), but the context of use is distinct too. This follows as the context of use for the [ANNOUNCEMENT OF NEW LEGAL STATUS] in (12b) can only be successfully deployed by suitably qualified participants. For instance, Neville Chamberlain was able to successfully deploy this lexical concept because on September 3rd 1939 when he declared war, he was the legally-appointed Prime Minister of The United Kingdom, and under the terms of The Royal Prerogative – powers invested in the monarch and deployed by the Prime Minister on behalf of the monarch – he was legally entitled to take the country to war.

Finally, the lexical concept which sanctions the use of *declared* in (12c) relates to the [ANNOUNCEMENT OF DUTIABLE GOODS AT CUSTOMS] lexical concept. This is distinct both in terms of illocutionary point and context(s) of use from the previously mentioned lexical concepts. This lexical concept is specialised for use in contexts involving customs provision at international border crossings. Its communicative function has to do with signalling as to goods being transported, or caused to be transported by the person issuing the ‘declaration’ in this specific context, with respect to restrictions on the nature and/or amount of goods that may be transported into the country which establishes the customs provision, and/or tax payable on particular goods.

Based on the foregoing discussion, Table 2 presents a summary of the key components of pragmatic point that are encoded as part of the linguistic content of each of the three lexical concepts. Much of the content associated with the three lexical concepts for *declared* comes from the conceptual content to which they afford access. However, pragmatic point, which concerns linguistic content, is highly schematic in nature. In these terms then, the distinction between the three lexical concepts relates to whether they stipulate that the setting is restricted or not, whether the participants are restricted or not, and the nature of the communicative function: the illocutionary point. Hence, by way of illustration, the lexical concept [ANNOUNCEMENT OF NEW LEGAL STATUS] encodes

linguistic content involves only the most generic aspects, including schematic information concerning the types of context in which a particular lexical concept can be deployed, the nature of the participants involved and the conditions which must hold.

Table 2. Pragmatic point for three lexical concepts of declared

Lexical concept	Setting	Participant(s)	Illocutionary point
[FORTHRIGHT INFORMATIONAL ASSERTION]	Unrestricted	Unrestricted	Make statement
[ANNOUNCEMENT OF NEW LEGAL STATUS]	Unrestricted	Restricted	Change official state
[ANNOUNCEMENT OF DUTIABLE GOODS AT CUSTOMS]	Restricted	Restricted	Make official statement

the following: there is no restriction on where the utterance can take place for it to realise its illocutionary point; the participants involved are, however, restricted, and the communicative purpose is to change some institutional state. This information is clearly highly schematic. However, it adequately captures, I argue, the highly stable aspects of the content associated with this lexical concept, which is to say, its linguistic content.

5. The nature of conceptual structure

In this section I am concerned, in broad terms, with conceptual structure: the nature and organisation of concepts. In LCCM Theory conceptual structure is modelled in terms of the cognitive model. A cognitive model is, in essence, similar to Barsalou's (1999) notion of a *simulator*.⁸

A simulator (Barsalou 1999) constitutes records of perceptual states, stored in a coherent format referred to as a *frame*, which can be re-activated, often in novel ways. The re-activations are referred to as *simulations*, (e.g. Barsalou 1999, 2003; Gallese and Lakoff 2005; Kaschak and Glenberg 2000; Prinz 2002; Glenberg and Kaschak 2002; Zwaan 1999, 2004).

8. The use of a novel term, 'cognitive model', is done for two reasons. Firstly, at this stage in our understanding, it is not clear to what extent units of semantic structure: lexical concepts, facilitate access to the conceptual system. For instance, the common experience of 'not being able to put thoughts into words', particularly as applied to subjective experiences, suggests that the linguistic system may be less well connected to certain types of conceptual representations than others. Indeed, this is a point made by Jackendoff (e.g., 1992). It is conceivable that some aspects of conceptual structure may only be partially accessible or even inaccessible to the linguistic system. I introduce the theoretical construct of the cognitive model, then, to distinguish between those simulators which are accessible via linguistic representations, and those which are not. Simply put, while the conceptual system is populated by simulators (Barsalou 1999), cognitive models are simulators which are specialised for being accessed by lexical concepts. Hence, the rationale for introducing the term 'cognitive model' is to identify those simulators with which the linguistic system interacts. The second reason is as follows. In his development of the notion of a simulator, Barsalou is primarily focused on the perceptual basis – in the wider sense as described earlier – of conceptual structure. While he acknowledges that other forms of information are likely to feed into conceptual representations, he is primarily exercised by accounting for the perceptual grounding of cognition. In my account, I explicitly acknowledge that propositional (i.e., non-perceptual) information may also become incorporated in cognitive models, which supplements the perceptual information already present. Such propositional information is likely to accrue via linguistically mediated routes, including narrative, exchange of news, and gossip. For these reasons, it is useful to distinguish the theoretical construct under development here, by applying the novel term cognitive model.

The mechanism known as *simulation* represents a general purpose computation performed by the conceptual system in order to recover the bodily states stored within frames and to perform operations deploying such perceptual states. As such, a frame can give rise to a potentially limitless set of re-activations or simulations. Hence, a simulator, and thus a cognitive model, encompasses a frame and a potentially unlimited set of simulations.

Following Barsalou, and indeed others who take an *embodied* or *grounded cognition* perspective (see Barsalou 2008 for a review), I assume that the perceptual states that make up cognitive models derive from a number of sources, as briefly introduced earlier. These include: i) the processing of external stimuli via sensory (or modal) systems (vision, audition, olfaction, haptics, and gustation); ii) action, which provides motor information relating to bodily states via proprioception: information about movements involving joints and muscles, as well as the vestibular system, which provides information as to position in space and motion trajectories.⁹ In addition, subjective (or introspective) experiences are just as important for giving rise to records of perceptual states that make up cognitive models. For instance, Damasio (e.g. 1994) in ground-breaking work on emotion has emphasised a number of categories of feelings that arise from internal body states. These include body states (emotions) that we label as Happiness, Sadness, Anger, Fear, and Disgust. These give rise to phenomenologically real, in the sense of directly experienced, *feelings*. Damasio identifies a further category of feeling, what he terms *background feelings*, which derive from internal body states. Background feelings arise from, among other things, interoceptive experience, which is to say the visceral sense – our felt sense of the internal organs and other internal bodily states. Other subjective experiences, which are directly felt, include various aspects of temporal experience which arise from bodily states (circadian rhythms such as the wake-sleep cycle), as well as perceptual processing, which is subserved by a wide range of neurologically instantiated temporal mechanisms (see Evans 2004a, 2004b and references therein), and consciousness (Chafe 1994; Grady 1997). Hence, the perceptual states that make up the frames and give rise to the simulations that comprise given cognitive models are grounded in both sensory-motor experience and subjective experience: experience of internal bodily and cognitive states, including emotion, mood and affect.

5.1 Perceptual symbols

Following Barsalou (1999) I assume that individual records of perceptual states are stored as *perceptual symbols*. It is well known from research on attention that during perceptual experience the cognitive system can focus attention on individual components of the stimulus array. For instance, attention can selectively focus on the colour of an object, filtering out, for instance, its shape, or texture, and even the surrounding objects (Garner 1974, 1978). Through selective attention, individual perceptual components derived from perceptual experience of the kinds discussed above are recorded, in bottom-up fashion, in

9. See Evans (to appear b) for a review of the operation of the sensory mechanisms responsible for sense-perception.

sensory-motor areas of the brain (Barsalou 1999).¹⁰ The components are stored in schematic fashion. This means that it is not individual perceptual states that are stored, but rather commonalities are abstracted across specific instances of perceptual states providing individual memories deriving from sense perception (e.g., individual memories for *red*, *hot* and *purr*), proprioception (e.g., *lift*, *run*) and subjective experience (e.g., *compare*, *similar*, *hungry*). These schematic memories Barsalou refers to as perceptual symbols. They are symbols in the sense that, later, in top-down fashion, they can be reactivated, or simulated, and can be used to support the range of symbolic behaviours that subserve a fully functional conceptual system.

Perceptual symbols implement a conceptual system as follows. Barsalou argues that memories of similar and related components become organised into a system of perceptual symbols which exhibit coherence: a *perceptual symbol system* (Barsalou 1999). This perceptual symbol system is what I refer to as a frame. A frame, then, is an information structure consisting of large collections of perceptual symbols, encoding information which is stable over time as well as incorporating variability. Hence, a frame provides a unified, and hence coherent, representation of a particular entity. For instance, a frame involves numerous components that have a perceptual basis, that are related in various ways. In addition, the perceptual symbols that collectively comprise the frame can be combined in a range of ways, giving rise to an infinite variety of simulations. Hence, a system of perceptual symbols gives rise to both a frame: a relatively stable knowledge matrix, and dynamic simulations.

5.2 Frames

In this section I identify a number of frame types. I do so based on Barsalou's work on frames (e.g. Barsalou 1991, 1992), and Barsalou et al. (1993). In broad terms, frames can be identified which relate to *things* and to *situations*. Further, within each of these broad divisions there are frames which are *episodic*, relating to specific types of experience and/or knowledge and frames which are *generic*, relating to schematisation over broadly similar aspects of experience and/or knowledge. The distinct frames identified below are *individuals* (episodic) and *types* (generic), which relate to things, and *episodic situations* and *generic situations*, which relate, self-evidently, to situations. I begin by focusing on the frames for things: individuals and types, before proceeding with a discussion of the frames for situations.

The world model

Barsalou et al. (1993) provide an ontology for a theory of knowledge representation, which is based on what they refer to as the *world model*. This comprises a person's beliefs about

10. There is compelling neuropsychological and neuroimaging evidence which supports the view that human conceptual representations are grounded in the modalities, and hence are perceptual in nature. For instance, categorical knowledge is grounded in sensory-motor regions of the brain (for reviews see Damasio 1989; Pulvermüller 1999, 2001). Damage to a particular sensory-motor region serves to impair the processing of categories that use the region in question to perceive physical exemplars.

the current state of the world. These beliefs relate to individuals, their current states and where they are located. Barsalou suggests that people employ a hierarchically-arranged core of spatial frames. That is, people represent the world and its contents in a spatial fashion, corresponding to continents, countries, cities, neighbourhoods, individual buildings, rooms and locations within rooms. They further locate entities within these locations, and integrate the spatial frames with temporal knowledge, for instance, relating to cycles and time-frames of various sorts including the seasons, the calendar, and temporal intervals such as years, months, weeks and days, as well as content-based temporal structures such as knowledge relating to one's own and family members daily routine, development over the life span, stages in career progression, and so on. Temporal information serves to organise past, present and future information in the world model and, Barsalou argues, does so orthogonally to the spatial core. Moreover, in this world model, people represent people's interactions and movements, updating the model continuously. For instance, while at work, a person might represent their partner's movements, going to the shops, returning home, or their children's activities while at school, and so on. People also represent other ongoing activities taking place in the various regions represented in their world model. For instance, one might know about a meeting of a University Exam Board taking place in a committee room near one's office, it being Tuesday afternoon, Prime Minister's question time taking place at the House of Commons, knowing – based on having read today's newspaper – that the Queen is currently staying at Windsor Castle rather than Buckingham Palace, that Big Ben in London is currently undergoing repairs and hence not presently chiming, and so on.

In the world model, two distinct kinds of frames can be distinguished which relate to things: individuals and types (Barsalou et al. 1993).¹¹ Individuals are frames that relate to animate and inanimate entities that are held to persist continuously in the environment. As such, individuals are central to the ontology of the world model. Individuals provide relatively stable information about a given entity: information that is both stable over time, as well as incorporating episodic information. Hence, the new information for a given individual is added to the frame thereby updating it on an ongoing basis. An individual is updated based on encounters with the entity it represents. For instance, the frame for 'my car' might include the petrol gauge reading the last time I interacted with it, and the fact that I have noticed there is an oil leak, and that the car needs cleaning. This information is merged into the frame to provide an updated representation.

Crucially, although the same individual may be encountered in the world on many occasions, often in the same day, in terms of the world model all the episodic information extracted during these encounters is integrated into the individual frame. This follows from the *one-entity one-frame principle* (Barsalou et al. 1993). This principle holds that only one frame can relate to any given entity. Hence, all the information extracted from experience, which is related to a particular individual, is merged into the frame for that

11. Barsalou et al. (1993) use the term 'model' to refer to what I am here calling 'type'. I prefer the more intuitively accessible term 'type' and also seek to avoid any confusion with the construct of the cognitive model. Hence, I do not use the term 'model'.

entity. Hence, the frame for a particular colleague at work may include information relating to his location the last time I interacted with him, and so on.

In addition to individuals, Barsalou et al. (1993) argue that there is another frame type which inheres in the world model. This type of frame, which I refer to as: type, is an abstraction across frames for individuals providing a frame for a type of individual. As such, types are not conceptualised as having corresponding entities in the world. For example, while the individual for 'my car' in the world model corresponds to my car in the world, the frame for 'car' is a type, and relates to a type of individual, abstracted from across a range of individuals. Hence, people understand their frames for types to inhere only in the world model, but not, crucially, in the world itself.

One of the features of individuals in the world is that they change location. In the world model, this feature is captured in terms of the phenomenon referred to as *transcendence* (Barsalou et al. 1993). Transcendence has to do with the number and range of locations at which individuals and types are represented. For instance, a colleague from work will be represented at work. However, a chance meeting at the local supermarket will ensure that the individual frame for the colleague becomes additionally stored at the supermarket location in the world model. When the colleague goes on vacation to Paris, and sends a postcard in to the office to report on the vacation, the individual is additionally stored as part of the Paris location in the world model.

Barsalou et al. (1993) argue that transcendent frames for individuals and types, while being located at multiple sites in the world model, become *functionally detached* from the world model. That is, they give rise to a level of information about the nature of individuals and types, and the interactions they can engage in which become abstracted from the spatial frames that form the core of the world model. In other words, transcendence gives rise to de-contextualised representations which form *transcendent taxonomies*. For example, the type for 'heart' is a feature of all mammals. Hence, its presence as part of the frame for numerous individuals and types gives rise to transcendence.

This property serves two important functions. Firstly, transcendence provides an important means of organising beliefs about the nature of entities in the world. It does so as it serves to capture similarities between individuals and models. As such, it facilitates inferences. For instance, we can infer that lions have hearts on the basis of knowing that all mammals possess hearts. Secondly, transcendent taxonomies may constitute important building blocks in the construction of the world model. This follows as transcendent information can be inserted into frames for new individuals upon first encounter. For instance, on encountering an unfamiliar cat, information from the model for cats is retrieved and copied, in order to form the basis for the new individual in the world model. This process serves to minimise the amount of learning about new entities before they can be adequately represented.

Situations

Having briefly described the ontology for individuals and types, I now consider how situations are modelled. The basic insight is that in addition to individuals and types, humans additionally represent situations, there being two kinds of situation: episodic situations

and generic situations. The distinction between episodic and generic situations is orthogonal to the distinction between individuals and types.

According to this approach, situations are part of larger *events* – events are composed of situations – while being made up of discrete *images*. As with situations, events and images, as I use the terms, are mental representations. The notions event, situation and image are somewhat akin to the notions of scripts, scenes and states developed in Schank (1975, 1982), and Schank and Abelson (1977), with the difference being that events, situations and images are made up of perceptual symbols, and hence are perceptual and thus embodied in nature.

One of the key insights of this approach is that it takes a *situated cognition* perspective. That is, people's frames for individuals and types are *situated*, and *local* rather than being de-contextualised and universal. An individual or type is situated in the sense that it is represented in the situations in which it occurs. For instance, the individual frame for 'my sofa' is represented as being located in my living room. Hence, the frame for my sofa is related to the situation frame for 'my living room'. Similarly, individuals and types are local in the sense that they relate only to exemplars actually encountered, rather than being generalised to entities universally. For instance, the type for 'sofa' incorporates information relating only to sofas that have been encountered. In this way, this approach to knowledge representation assumes that the conceptual system is directly grounded in situated action and interaction.

Barsalou et al. (1993) propose that the mental representations they refer to as images are static spatial scenes (cf. Tyler and Evans 2003). These may consist of frames for individuals and/or types, viewed from a particular viewpoint, with a particular geometric, topological and functional relationship holding between them. Crucially, an image is composed of numerous perceptual symbols. For instance, a person may represent a picture hanging on the wall above the sofa in their living room.

A situation is comprised of a series of images. Hence, and as with an image, a situation may consist of a relatively stable set of individuals and types. The difference is that a situation, while occupying a relatively constant region of space is dynamic, in the sense that entities may interact and move around, and there is change over time. For instance, a situation might involve a person approaching the sofa, sitting down, turning their head to look at the picture on the wall, turn their head away again, sitting for a while, before getting up and moving away from the sofa.

An event comprises a series of two or more situations which are related in coherent fashion. The key difference between an event and a situation is that an event involves a significant outcome, often involving a change in regions of space and/or the individuals and/or types involved in the event. For instance, an event might involve a person going to a department store and purchasing a picture, bringing it home in their car, fetching a hammer and nail from the garage, selecting a spot on the wall above the sofa to hang the picture, knocking a nail in the wall at the desired location, and hanging the picture above the sofa. A table summarising the differences between image, situation and event *qua* mental representations is provided in Table 3.

As observed above, there are distinct sorts of frames relating to both episodic and generic situations, which parallels the distinction between individuals and types. An episodic

Table 3. Features of images, situations and events

Features of images	Features of situations	Features of events
i. a set of perceptual symbols	i. a series of images	i. a series of two or more situations
ii. represents individuals and/or types	ii. depicts a relatively constant set of individuals and/or types	ii. the situations are related in a coherent manner
iii. a static spatial configuration	iii. depicts some significant change over time	iii. the situations lead to a significant outcome
iv. viewed from a particular perspective	iv. occurs in a relatively constant region of space	

Table 4. Identification of commonalities in the formation of an abstract situation (after Barsalou et al. 1993)**Two situations are related when the following occur:**

- i. They share a common number of images.
- ii. They share common individuals and/or types.
- iii. The configuration of individuals/types in each similar image across situations is qualitatively the same.
- iv. The transformations of individuals/types between similar images across situations is qualitatively the same.
- v. The two situations culminate in a common end state.

situation arises from perceiving a situation in the world, the situation *qua* frame constituting a mental representation of the perceived situation. Moreover, humans represent situations at the locations in their world model where the situation occurs. For instance, in the example of the situation involving the hanging of a picture above the sofa, the frame for the episodic situation is linked to the frame for the conceptualiser's living room. On this account, and just as we saw with frames for individuals above, episodic situations are not wholly episodic. They also include a potentially large amount of generic information. This is due to the phenomenon of transcendence, which facilitates cognitive economy: generic knowledge can be shared between related frames. As with frames for things: individuals and types, discussed above, frames for situations are associated with temporal knowledge structures such as those relating to daily routines, life periods, hours of the day, and so on.

In contrast, frames for generic situations do not include episodic information. Rather they develop by virtue of abstracting away points of difference, in order to distil the commonalities that persist in different frames for episodic situations. Like frames for types, discussed above, generic situations do not have direct counterparts in the world.

Barsalou et al. (1993) propose that frames for a generic situation are formed when two or more episodic situations share a number of commonalities. These are presented in Table 4. These commonalities serve to indicate that two episodic situations are related. The episodic situations in question are then abstracted in order to form a generic situation for this type of situation.

5.3 The structure of frames

Frames have three basic constituents: *attribute-value sets*, *structural invariants* and *constraints*. In this section, which draws on Barsalou (1992) I examine each of these in turn.

Attribute-value sets

Frames consist of sets of attributes and values. An attribute concerns some aspect of a given frame, while a value is the specification of that aspect. For example, in terms of the vastly simplified frame for CAR depicted in Figure 4, ENGINE represents one aspect of the CAR, as do DRIVER, FUEL, TRANSMISSION and WHEELS. An attribute is therefore a concept that represents one aspect of a larger whole. Attributes are represented in Figure 4 as ovals. Values are subordinate concepts, which represent subtypes of an attribute. For instance, SUE and MIKE are types of DRIVER; PETROL and DIESEL are types of FUEL; MANUAL and AUTOMATIC are types of TRANSMISSION, and so on. Values are represented as dotted rectangles in Figure 4. Crucially, while values are more specific than attributes, a value can also be an attribute, because it can also have subtypes. For instance, PETROL is an attribute to the more specific concepts UNLEADED PETROL and LEADED PETROL, which are values of PETROL. Attributes and values are therefore super-ordinate and subordinate concepts within an *attribute taxonomy*: subordinate concepts, or values, which are more specific, inherit properties from the super-ordinate concepts, or attributes, which are more general.

In addition, attributes within a frame can be associated with their own *attribute frame*, providing an embedded form of framing. For instance, the attribute DRIVER in the CAR frame may have a number of attributes associated with it, including AGE, SEX, STATUS OF

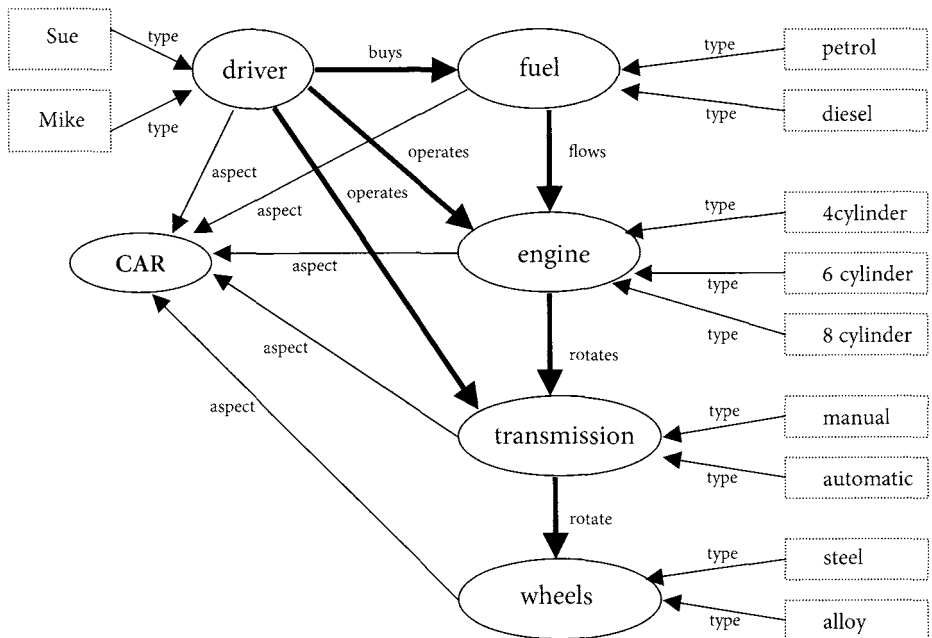


Figure 4. Frame for car (adapted from Barsalou 1992: 30)

DRIVING LICENCE (i.e. whether it is ‘clean’ or not), NUMBER OF YEARS EXPERIENCE, and so on. As frames are dynamic entities, undergoing continuous updating, attributes can be added to frames based on new encounters, or in order to achieve a particular goal. For instance, in the light of the recent introduction of a new banding scheme for road tax – an annual tax paid on all vehicles in the UK to use the public highway – based on petrol consumption, UK car owners are likely to have added a new attribute to their frame type for CAR relating to CAR TAX LEVEL. It is also worth emphasising that attribute-value sets, as with other aspects of knowledge representation, are likely to be idiosyncratic, and hence to vary from person to person.

A final property of attribute-value sets that I mention relates to what Barsalou and Billman (1989) have referred to as *attribute systematicity*. This concerns the idea that certain attributes are core, in the sense that they frequently recur across contexts.¹² This can facilitate frame formation. For instance, if a particular value for an attribute is not known when setting up a new frame of the type individual, a value for a core attribute can be ascribed based on the core attribute set retrieved from memory. For instance, imagine your friend is proudly showing off his new bright red sports car to you, a core attribute of the type frame: SPORTS CAR, is FUEL with the value PETROL. Hence, even though there may be no direct evidence that the car takes fuel, for instance, because you haven’t noticed a petrol cap, or seen evidence of a fuel tank, this is something that will be added to the frame for this individual, and the value PETROL will be added as a consequence.

Structural invariants

According to Barsalou, “[A]ttributes in a frame are not independent slots but are often related correlationally and conceptually” (Barsalou 1992: 35). In other words, attributes within a frame are related to one another in consistent ways across *exemplars*: instances of a given frame in the world. For example, in most exemplars of the frame CAR it is the driver who controls the speed of the ENGINE. This relation holds across most instances of cars, irrespective of the values involved, and is therefore represented in the frame as a *structural invariant*: a more or less invariant relation between attributes DRIVER and ENGINE. In Figure 4 structural invariants are indicated by bold arrows. Hence, a structural invariant constitutes what Barsalou (1992) terms ‘a normative truth’ holding between attributes within a frame.

Constraints and factors

Like structural invariants, constraints and factors are relations that hold between attributes, or more specifically, between attribute values. However, rather than capturing normative relations, constraints and factors give rise to variability in the values associated with attributes. This follows as values in a given frame are interdependent on the values associated with other attributes. There are two kinds of constraints, which I briefly review

12. As is well known, correlations in experience give rise to associative strength in memory: co-occurrence gives rise to a core set of attributes, which thus exhibit systematicity. See references in Barsalou et al. (1993) for instance.

below, and two factors. The constraints are *global constraints*, and *local constraints*. The two factors are *contextual factors* and *goal factors*.

Global constraints serve to constrain attribute values globally. This means that a modification in one value entails a proportional modification in a related value. For instance, consider the example of a TRANSPORTATION frame involving a passenger in a taxi, for instance, being transported from one location to another. In this frame there is a negative attribute constraint which holds between the attributes SPEED and DURATION. That is, as the value for the attribute SPEED increases (and transportation becomes faster), so the value for the attribute DURATION decreases.

Local constraints constrain sets of values locally, rather than globally. That is, the presence of a given value entails the presence of a related value, while the absence of one entails the absence of another. For instance, consider a frame for VACATION. If the attribute ACTIVITY has the value SKIING, then this requires that the attribute HOLIDAY DESTINATION has the value SKI RESORT. Similarly, if the attribute ACTIVITY has the value SURFING, then the destination attribute must have the value OCEAN BEACH.

Contextual factors relate to aspects of context which serve to influence attribute values. For instance, the activity of SKIING requires a SKI RESORT, while increasing SPEED of travel reduces the DURATION of the journey. As aspects of situations are related rather than being independent, context constitutes a factor which can influence both global and local constraints.

Now I consider goal factors. In addition to context, an agent's goal(s) also provides a factor that influences the interaction between values associated with related attributes. For instance, in a PHYSICAL WORK OUT frame, the agent's goal, to get fit, serves to ensure that the attribute EXERTION forms part of the frame.

5.4 Chaining within the conceptual system

In this section I briefly consider the phenomenon of *chaining* (Barsalou et al. 1993; see also Lakoff 1987). The conceptual system is not a haphazard collection of cognitive models. Rather, cognitive models exhibit a range of often complex interconnections. As such, cognitive models are linked in a web of interconnections, of diverse sorts: hence, chaining. The consequence of this, in terms of linguistic interaction, is that access sites established by lexical concepts provide a deep semantic potential for purposes of linguistically-mediated communication.

Chaining is a consequence of a number of different types of interconnections and relationships holding between frames. One such interconnection arises due to the phenomenon of attribute frames, discussed above. That is, frames are embedded within larger frames. Take the frame CAR, discussed above. A salient attribute associated with this type is ENGINE. The knowledge of engines possessed by one group of human conceptualisers, namely car mechanics, is highly complex, and this attribute includes many subordinate attributes each with corresponding values, which are themselves subordinate attributes with further values, and so on. In this way, a frame subsumes multiples frames which are embedded, capturing aspects of the larger units of which they are subparts.

Another way in which chaining occurs arises from the phenomenon of transcendence. This relates to the situated nature of cognitive models for things: individuals and types. Recall that cognitive models of this kind are 'located' in situations. In other words, cognitive models for things are located in the world model at the points at which they are encountered. Hence, cognitive models for episodic and generic situations include representations for individuals and types. The greater the number of situations to which individuals and types are linked the greater their transcendence is held to be. Hence, transcendence is a function of how interconnected cognitive models for things are with the range of representations for situations, and hence events, with which they are connected.

Another motivation for chaining arises due to the componential nature of the conceptual system itself. Recall that cognitive models are comprised of sets of perceptual symbols. As perceptual symbols are records of discrete perceptual states (e.g. *purr*, *red*, *hot*, etc.), similar perceptual symbols (e.g. *red*) form part of many different cognitive models within the conceptual system. As such, unique records of similar perceptual states persist throughout the conceptual system. The consequence of this is that the conceptual system is thorough-goingly redundant in terms of the nature of the representations which make up the range of cognitive models which populate it. This provides, naturally, commonalities across cognitive models, and is a consequence of a fundamental design feature of the conceptual system.

Another way in which chaining arises is due to the relationships that exist between cognitive models, due to, broadly, the distinction between episodic versus generic cognitive models. For instance, in terms of cognitive models for things, we have the distinction individuals and types. While individuals may be related to each other based on the dimensions of chaining mentioned in the preceding paragraphs, a type is related to all the individuals from which it is formed. Similarly, a generic situation is related to all the episodic situations that it resembles, and from which it has abstracted across to provide a generic situation.

6. Interaction between the linguistic and conceptual systems

A key feature of knowledge representation in humans is that the linguistic system interacts with the conceptual system in order to facilitate access to conceptual knowledge. Indeed, as the philosopher of science Jesse Prinz (2002: 14) has observed:

Concepts must be capable of being shared by different individuals and by one individual at different times. This requirement...must be satisfied if concepts are to play some of their most important explanatory roles...it is almost universally assumed that concepts play a pivotal role in linguistic communication.

Indeed, a fundamental design feature of human cognition is that linguistic representations provide an indexing and control function, greatly increasing the range of uses and flexibility of the human conceptual system. However, this does mean that linguistic representations are equivalent to the concepts which populate the conceptual system.

I assume that the human conceptual system is, *en grandes lignes*, essentially the same as the primate conceptual system. Recent findings suggest that such an assumption is not unreasonable (e.g. Barsalou 2005; Hurford 2007). Given the relatively recent emergence of language, and the far greater antiquity of the conceptual system¹³ I assume that linguistic representations evolved to complement and enhance the existing form of representations that inhere in the conceptual system, rather than duplicating them.

From the perspective of LCCM Theory, the interaction between the linguistic and conceptual systems is facilitated by what I earlier referred to as open-class lexical concepts. I discuss the nature of the interaction by examining some of the relevant issues below.

6.1 Access sites

The primary way in which the representations inhering in the linguistic and conceptual systems interact is by virtue of *access sites*. An access site is a theoretical construct in LCCM Theory which represents a composite of the range of *association areas* that hold between an open-class lexical concept and the conceptual system. An association area is a location in the conceptual system with which a specific lexical concept is associated. In other words, an association area provides a point of convergence between the two systems facilitating interaction between content from both. As a given lexical concept has typically many association areas, an access site constitutes the set of association areas for a given lexical concept. For example, and as we shall see below, the lexical concept [RED] is associated with many representations for individuals and types, each with its own distinctive hue, throughout the conceptual system. All the association areas collectively comprise the access site for this lexical concept. Yet the complexity of the way in which [RED] facilitates access to conceptual structure gives rise, as we shall see, to a large *semantic potential*.

The purpose of an access site is to facilitate integration of linguistic and conceptual content in order to provide an *integrated simulation*.¹⁴ Hence, the evolutionary motivation, on this account, for the linguistic and conceptual systems to interact is in order to make use of conceptual structure inhering in the conceptual system in service of linguistically-mediated communication.

I hypothesise that the association areas that comprise an access site arise by virtue of usage patterns: vehicles sanctioned by specific lexical concepts being used in the context of perceived things and situations. Based on such patterns of use, statistical frequencies are extracted which serve to associate lexical concepts with the regions of the conceptual system where such things and situations are represented, giving rise to association areas. Access sites are thus probabilistic, in the sense that the greater the frequency with which

13. For discussion, a sample of relevant book length treatments from various perspectives include Corballis (2002), Deacon (1997), Donald (1991), Dunbar (1996), Mithen (1996), Hurford (2007). See also the excellent collection of papers in Christiansen and Kirby (2003).

14. An integrated simulation is equivalent to what I what I have referred to as a conception: the meaning derived from compositional processes involved in understanding a well-formed utterance.

a language user experiences a sanctioning lexical concept and a thing/situation as co-occurring, the greater the strength of the association area.¹⁵

6.2 Semantic potential

One consequence of the chaining exhibited by the conceptual system is that lexical concepts, by encoding access sites, facilitate access to a large semantic potential. To illustrate, let's briefly re-consider the lexical concept [RED] associated with the form *red*. The lexical concept [RED] facilitates access to a bewildering number of distinct perceptual symbols which contribute to a vast number of cognitive models in the conceptual system of any language user of English. To get a sense of the semantic potential involved, consider all the individuals and types that a single person will represent in their world model that features the perceptual state I gloss as *red*.

Limiting ourselves to types we might list Royal Mail post boxes, red squirrels, foxes, roses, blood, lipstick, Santa Claus' clothes, a robin's throat, strawberries, the red stop sign on the public highway, tomatoes, red traffic light, red cross, the flag of St. George, celebrity carpets, Babybel cheese wax, chilli peppers, fire engines, the Chinese flag, red wine, superman's cape, fire, henna, and so on. Notice that the represented hue associated with these types may vary from person to person, based on cultural experience, and so on. Nevertheless, we can imagine contexts in which we would apply the phonological vehicle *red* in order to evoke the colour associated with these types.

In addition, there are further situations, both episodic and generic, that involve the individuals and types which include a perceptual symbol that I gloss as *red*. However, each of these perceptual symbols is unique to the individual and/or type and hence the situation of which it forms part. After all, it is the generic situation in which a teacher scrawls red ink on a pupil's exercise book, evoking a different perceptual symbol than the one evoked when we simulate a red squirrel scurrying up a tree. Nevertheless, the lexical concept [RED] is associated with, and hence facilitates access to, both. Put another way, the semantic potential for the lexical concept [RED] comes from the diverse range of perceptual symbols that are found in these cognitive models, and many others. Moreover, it is precisely because [RED] facilitates access to such a diverse potential that the vehicle *red* exhibits such variation in the way it can be used, as exhibited by the very different simulations we achieve for 'red' in the examples discussed earlier in Section 3: the red associated with a red squirrel versus the red ink of a school teacher's pen.

6.3 The uniqueness of the access site

While lexical concepts are typically associated with a number (often many) cognitive models, which thereby make up the access site, the exact nature of the access site with which a lexical concept is associated is held to be unique. Put another way, no two lexical concepts

15. See Barsalou et al. (To appear) for discussion of a related proposal. See also Boroditsky and Prinz (to appear).

share the same access site. While the range of cognitive models to which lexical concepts may be similar, they may never be exactly the same. The consequence of this is that each lexical concept has a unique *cognitive model profile*: the range of cognitive models which make up an access site. From the perspective of the linguistic system, this means that there can be no true synonymy between lexical concepts.

To illustrate, consider the lexical concepts which I gloss as [SHORE] and [COAST] associated with the forms *shore* and *coast*, respectively. As observed by Fillmore (1982) while the semantic representation for these two lexical concepts is similar it is not identical. This follows, in present terms, as while each of these lexical concepts exhibits partial overlap in the primary cognitive models, there are also distinctions. For instance, both lexical concepts facilitate access to a cognitive model profile relating to the strip of land that borders land and sea. However, each lexical concept accesses a cognitive model relating to a generic situation from which this land region is viewed. In the case of [SHORE] this concerns a sea-based perspective, i.e., on board a ship. In contrast, [COAST] does so from the perspective of land-based location. For this reason, a shore-to-shore trip is across water while a coast-to-coast trip is over land.

7. Summary

This paper has been concerned with developing an account of semantic representation, as assumed by LCCM Theory. LCCM Theory assumes a principled separation between the evolutionarily earlier conceptual system and the more recent linguistic system. Each system is populated by different types of 'semantic' representation: the lexical concept and the cognitive model. Moreover, the nature of the content associated with the two systems is of a fundamentally different type. Linguistic content, encoded by lexical concepts, is highly schematic in nature, providing a structuring function to simulations. In contrast, conceptual content, encoded by cognitive models, provides perceptually rich and highly detailed information. In addition to encoding linguistic content, a subset of lexical concepts – open-class lexical concepts – serve as access sites, thereby facilitating interaction between linguistic and conceptual content, thereby giving rise to integrated simulations. In essence, LCCM Theory assumes that the linguistic system provides an executive control function, allowing access to conceptual representations for purposes of linguistically-mediated communication.

References

- Barsalou, L. 1991. Deriving categories to achieve goals. In *The Psychology of Learning and Motivation: Advances in Research and Theory*, Volume 27, G. H. Bower (ed.), 1–64. San Diego, CA: Academic Press.
- Barsalou, L. 1992. Frames, concepts, and conceptual fields. In *Frames, Fields, and Contrasts: New Essays in Lexical and Semantic Organization*, A. Lehrer & E. F. Kittay (eds), 21–74. Hillsdale, NJ: Lawrence Erlbaum.
- Barsalou, L. 1999. Perceptual symbol systems. *Behavioral and Brain Sciences* 22: 577–660.

- Barsalou, L. 2003. Situated simulation in the human conceptual system. *Language and Cognitive Processes* 18: 513–562.
- Barsalou, L. 2005. Continuity of the conceptual system across species. *Trends in Cognitive Sciences* 9: 309–311.
- Barsalou, L. 2008. Grounded cognition. *Annual Review of Psychology* 59: 617–645.
- Barsalou, L. & Billman, D. 1989. Systematicity and semantic ambiguity. In *Resolving Semantic Ambiguity*, D. Gorfein (ed.), 146–203. New York: Springer-Verlag.
- Barsalou, L., Yeh, W., Luka, B., Olseth, K., Mix, K. & Wu, L.-L. 1993. Concepts and meaning. In *Chicago Linguistics Society 9: Papers from the Parasessions on Conceptual Representations*, Volume 2, K. Beals, G. Cooke, D. Kathman, K. E. McCullough, S. Kita & D. Testen (eds), 23–61. Chicago Linguistics Society.
- Barsalou, L., Santos, A., Simmons, K. & Wilson, C. D. To appear. Language and simulation in conceptual processing. In *Symbols, Embodiment, and Meaning*, M. De Vega, A. Glenberg & A. Graesser (eds). Oxford: Oxford University Press.
- Boroditsky, L. & Prinz, J. To appear. What thoughts are made of. In *Embodied Grounding: Social, Cognitive, Affective, and Neuroscientific Approaches*, G. Semin & E. Smith (eds). Cambridge: Cambridge University Press.
- Chafe, W. 1994. *Discourse, Consciousness, and Time: The Flow and Displacement of Conscious Experience in Speaking and Writing*. Chicago: University of Chicago Press.
- Christiansen, M. H. & Kirby, S. 2003. *Language Evolution*. Oxford: Oxford University Press.
- Corballis, M. 2002. *From Hand to Mouth: The Origins of Language*. Princeton: Princeton University Press.
- Damasio, A. 1989. Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition* 33: 25–62.
- Damasio, A. 1994. *Descartes' Error: Emotion, Reason and the Human Brain*. London: Vintage.
- Deacon, T. 1997. *The Symbolic Species: The Co-Evolution of Language and the Brain*. New York, NY: W.W. Norton and Company.
- Donald, M. 1991. *Origins of the Modern Mind: Three Stages in the Evolution of Culture and Cognition*. Harvard: Harvard University Press.
- Dunbar, R. 1996. *Grooming, Gossip and the Evolution of Language*. London: Faber and Faber.
- Evans, V. 2004a. *The Structure of Time: Language, Meaning and Cognition*. Amsterdam: John Benjamins.
- Evans, V. 2004b. How we conceptualise time. *Essays in Arts and Sciences XXXIII* (2): 13–44.
- Evans, V. To appear a. *How Words Mean*. Oxford: Oxford University Press.
- Evans, V. To appear b. The perceptual basis of spatial representation. In V. Evans & P. Chilton (eds). London: Equinox.
- Fillmore, C. 1982. Frame semantics. In *Linguistics in the Morning Calm*, 111–137. Seoul, South Korea: Hanshin Publishing Co.
- Fillmore, C. 1997. *Lectures on Deixis*. Stanford: CA: CSLI.
- Fillmore, C., Kay, P. & O'Connor, M. C. 1988. Regularity and idiomatity in grammatical constructions: The case of *let alone*. *Language* 64: 501–538.
- Gallese, V. & Lakoff, G. 2005. The brain's concepts: The role of the sensory-motor system in reason and language. *Cognitive Neuropsychology* 22: 455–479.
- Garner, W. R. 1974. *The Processing of Information and Structure*. Hillsdale, NJ: Lawrence Erlbaum.
- Garner, W. R. 1978. Selective attention to attributes and to stimuli. *Journal of Experimental Psychology: General* 107(3): 287–308.
- Glenberg, A. & Kaschak, M. 2002. Grounding language in action. *Psychonomic Bulletin and Review* 9: 558–565.
- Goldberg, A. 1995. *Constructions. A Construction Grammar Approach to Argument Structure*. Chicago: University of Chicago Press.

- Grady, J. 1997. *Foundations of Meaning: Primary Metaphors and Primary Scenes*. PhD Dissertation, University of California at Berkeley.
- Hurford, J. 2007. *Origins of Meaning*. Oxford: Oxford University Press.
- Jackendoff, R. 1992. *Languages of the Mind: Essays on Mental Representation*. Cambridge, MA: MIT Press.
- Kaschak, M. & Glenberg, A. 2000. Constructing meaning: The role of affordances and grammatical constructions in sentence comprehension. *Journal of Memory and Language* 43: 508–529.
- Lakoff, G. 1987. *Women, Fire and Dangerous Things: What Categories Reveal about the Mind*. Chicago: Chicago University Press.
- Langacker, R. 1987. *Foundations of Cognitive Grammar: Volume 1 Theoretical Prerequisites*. Stanford: Stanford University Press.
- Levinson, S. C. 1983. *Pragmatics*. Cambridge: Cambridge University Press.
- Mithen, S. 1996. *The Prehistory of the Mind: A Search for the Origins of Art, Religion and Science*. London: Orion Books.
- Prinz, J. 2002. *Furnishing the Mind: Concepts and their Perceptual Basis*. Cambridge, MA: MIT Press.
- Pulvermüller, F. 1999. Words in the brain's language. *Behavioral and Brain Sciences* 22: 253–336.
- Pulvermüller, F. 2001. *The Neuroscience of Language: On Brain Circuits of Words and Serial Order*. Cambridge: Cambridge University Press.
- Schank, R. 1975. *Conceptual Information Processing*. New York: Elsevier.
- Schank, R. 1982. *Dynamic Memory: A Theory of Reminding and Learning in Computers and People*. Cambridge: Cambridge University Press.
- Schank, R. & Abelson, R. 1977. *Scripts, Plans, Goals, and Understanding*. Hillsdale, NJ: Lawrence Erlbaum.
- Searle, J. 1969. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge: Cambridge University Press.
- Talmy, L. 2000. *Toward a Cognitive Semantics Volume 1*. Cambridge, MA: MIT Press.
- Tyler, A. & Evans, V. 2003. The Semantics of English Prepositions: Spatial Scenes, Embodied Meaning and Cognition. Cambridge: Cambridge University Press.
- Zwaan, R. 1999. Embodied cognition, perceptual symbols, and situation models. *Discourse Processes* 28: 81–88.
- Zwaan, R. 2004. The immersed experienter: toward an embodied theory of language comprehension. In *The Psychology of Learning and Motivation*, B. H. Ross (ed.), 35–62. New York, NY: Academic Press.

Behavioral profiles

A corpus-based approach to cognitive semantic analysis

Stefan Th. Gries and Dagmar Divjak*

1. Introduction

In this paper we will look into questions that concern what may be considered two of the central meaning relations in semantics, i.e. polysemy or the association of multiple meanings with one form and synonymy, i.e. the association of one meaning with multiple forms.

In the domain of polysemy, cognitive semanticists typically face issues which center on the questions of how to determine whether two usage events are sufficiently similar to be considered instantiations of a single sense and how to establish the prototypicality of a sense/several senses; we adopt Evans's (2005: 33, n. 2) definition of sense as those meanings which have achieved conventionalization and are instantiated in semantic memory. In the domain of near synonymy, semanticists need to uncover among other things what syntactic, semantic and/or pragmatic differences there are between near synonyms and what the semantic and/or functional relation is between near synonyms in a semantic space. In order to solve these problems they need to be able to measure the degree of similarity between senses and/or words and to decide how and where to connect a sense/word to another sense/word in a network.

Several solutions to these problems have been put forward in the literature, in particular for polysemy-related issues. One such solution for polysemy-related issues is the full-specification approach inspired by Lakoff and his collaborators (cf. e.g. Norvig and Lakoff 1987; Lakoff 1987) where minimal perceived differences between usage events constitute different senses and image schemas. Related to this is Kreitzer's (1997) partial-specification approach where information from three different levels of schematization – the so-called component, relational, and integrative levels – is integrated, yet minimally different usage events need not constitute different senses. Both of these approaches suffer from methodological inadequacies and representational problems, however. As for the former approach, information provided by the context the word under study occurs in is not taken into account (cf. Sandra and Rice 1995; Tyler and Evans 2001), there is no

* At the time of writing, Dagmar Divjak was a Postdoctoral Research Fellow of the FWO-Vlaanderen (Belgium), working in the Research Unit QLVL at the KU Leuven (Belgium). The financial support of the FWO is gratefully acknowledged.

method for identifying how the primary sense has developed, and empirical support for fine-grained semantic distinctions is not provided. As for the latter approach, problems relate to the vagueness of the representation and the lack of clarity concerning the status of the proposed networks.

Contrary to the above-mentioned studies, Sandra and Rice (1995) and Rice (1996) measure the similarity of senses using a variety of experimental methods such as off-line sentence sorting followed by hierarchical cluster analyses, off-line sentence similarity judgments, on-line acceptability judgments and sentence generation. While this experimental approach is certainly more objective than introspective approaches, it is also a bit problematic. First, it remains unclear to what degree sentential context rather than the prepositions under investigation influence the subjects' sorting style (cf. Klein and Murphy 2001, 2002) as does the influence of methodological choices on the clustering. Second, the questions remain whether subjects use the same cognitive strategies for conscious off-line classification as for subconscious on-line production (a general problem of experimental approaches) and whether conscious off-line classification reflects the patterns underlying mental representation.

More recently the principled-polysemy approach was introduced by Tyler and Evans (2001). Tyler and Evans argue that previous research on polysemy lacks a constrained approach to distinguishing senses. For example, in their work on *over* they propose that a distinct sense of *over* should be posited if and only if the meaning of *over* in one utterance (i) involves a different spatial configuration from *over*'s use in another utterance and (ii) cannot be inferred from encyclopedic knowledge and/or context. In later work (on *time*) within the same framework, Evans (2005: 41) introduces three criteria, which we quote here in detail because we will return to them later:

- i. a **meaning** criterion: a distinct sense must contain additional meaning compared to other already established senses;
- ii. a **concept elaboration** criterion: a distinct sense will feature unique or highly distinct patterns of concept elaboration [...] as in the lexical choices signaled by patterns of modification [...] or in the verb phrase which complements the noun phrase [...]. I assume that syntagmatic relations of this kind follow from semantic/conceptual considerations (see Croft's 2001 discussion of what he terms collocational dependencies);
- iii. a **grammatical** criterion: a distinct sense "may manifest unique or highly distinct structural dependencies. That is, it may occur in unique grammatical constructions".

Although the last two criteria are in fact predictions about distributional patterns of the words under study, so far the proponents of the principled-polysemy approach have not utilized corpus data.

The second major question we raised above, namely how to determine the prototypical sense(s) of a word, has been an issue in polysemy ever since the first cognitive-linguistic analyses appeared. A variety of criteria has been proposed to isolate the prototypical sense (cf. e.g. Rice 1996: 145–146; Tyler and Evans 2001: Section 3.3; Evans 2005: Section 2.2.3) and the following is a non-exhaustive list of such criteria: asymmetrical judgments of goodness or similarity; ease of elicitation; gradation within the category; diachronically earliest sense; centrality/predominance in the semantic network; use in composite forms;

frequency of occurrence etc. Unfortunately, it remains unclear whether all criteria can be applied to all kinds of words and sometimes the proposed criteria make conflicting or counter-intuitive predictions (cf. Corston-Oliver 2001; Divjak and Gries 2006; Gries 2006). We admit, though, that this is a risk of all multifactorial approaches rather than a problem of any one particular study mentioned above.

Although near synonymy constitutes, in a sense, the opposite of polysemy, it has received relatively little attention in recent years. Within cognitive linguistics but a few studies have been devoted to the phenomenon (Geeraerts 1985; Mondry and Taylor 1992; Taylor 2003); this is likewise the case within western linguistics in general (Cruse 1986 being the exception). Surprisingly, the studies that have been carried out within the cognitive linguistic framework do utilize non-elicited material, yet the illustrative use of the corpus data makes them but mere forerunners of the corpus-based approach we will introduce below (see also Divjak 2004).

To sum up, in spite of the prominence the term 'usage-based' currently enjoys in cognitive-linguistic publications and in spite of the fact that some approaches explicitly couch their criteria in corpus-linguistic terms, there are few truly corpus-based approaches to polysemy and near synonymy. One laudable exception is the largely corpus-based approach of Kishner and Gibbs (1996) to *just* (as well as Gibbs and Matlock 2001 on *make*) which anticipated much of the above mentioned proposals by Evans (2005). Gibbs and colleagues investigate R1 collocates and colligations, correlating different senses with collocations and colligations.¹ Their "findings suggest the need to incorporate information about [...] lexico-grammatical constructions in drawing links between different senses of a polysemous word" (Gibbs and Matlock 2001:234). Unfortunately, these studies do not fully utilize the potential of corpus data: citations in corpus data have more to offer than just individual collocations and colligations, and restricting the analysis to R1 collocates is a heuristic that is blind to syntactic structure (cf. points of critique also raised in collocation analysis; cf. Stefanowitsch and Gries 2003; cf. Divjak 2006).

Work in corpus linguistics, on the other hand, has exploited the potential of corpus data more fully. Such studies start out from the self-evident statement that corpus data provide distributional frequencies. The assumption then is that distributional similarity reflects, or is indicative of, functional similarity, the understanding of functional similarity being rather broad, i.e. encompassing semantic, discourse-pragmatic, and other functions a particular expression can take on. Against this background, Atkins's (1987) study on *danger* involves collocate analysis from L7 to R7, colligations, part of speech (POS) characteristics of the head word, and all the collocations/colligations correlating (probabilistically or perfectly) with a particular sense are referred to as an ID tag. Also, Hanks's paper (1996) on *urge* involves collocate and colligation analysis. He argues that "the semantics

1. The term *collocation* encompasses both the probabilistic co-occurrence of word forms (e.g. *different to vs different than*) as well as the absolute frozenness of expressions (e.g. *by and large*). *Collocations* are thus co-occurrences of words which are referred to as *collocates*; often, the letters L (for left) and R (for right) are used together with a number to refer to the position of one collocate with respect to the head word (e.g. R1 meaning 'the first collocate to the right'). The term *colligations* refers to the co-occurrence of word forms with grammatical phenomena (e.g. the preference of *consequence* to occur as a complement and with an indefinite article).

of a verb are determined by the totality of its complementation patterns” (1996: 77), where a set of coarse complementation patterns and semantic roles of a word is referred to as a behavioral profile. Unfortunately, neither Atkins nor Hanks provides conclusive evidence concerning the predictive power of the ID tags investigated. In addition, much of the method of analysis remains to be fleshed out and lacks quantitative sophistication.

In other words, while interesting studies have been conducted, semantic analyses in the area of polysemy and near-synonymy have often been based on introspective data. This makes them not only empirically problematic, but it likewise prevents the development of a rigorous, quantifiable, and objectively comparable methodology. Corpus-based or computational-linguistic studies, on the other hand, do introduce methodological rigor, yet, they are rather limited as they treat words with few different senses or focus on small sets of semantically similar words (*almost vs. nearly, high vs. tall, between vs. through*). In addition, they use data that constitute impoverished subsets of what is actually available: basing a semantic analysis of words solely on collocates in one sequentially defined slot means both seriously limiting the data taken into consideration and disregarding the syntactic structure of the clause under investigation. Thirdly, the databases used in computational linguistic research may be noisy or skewed given that such studies often rely on (semi-) automatic preprocessing tools.

In this paper, we will argue in favor of a radically corpus-based approach to polysemy and near synonymy. The approach is *radically* corpus-based because we rely on the correlation between distributional patterns and functional characteristics to a much larger extent than most previous cognitive-linguistic work; we will clarify this statement below. We submit our approach is a worthwhile addition to the cognitive-semantic field: the notion *usage-based* is encountered more and more frequently – the principled-polysemy approach even makes explicit use of corpus-linguistic terms – and corpus-based approaches have a variety of advantages that include, but are not limited to, the following:

1. the criteria used are not based on traditional minimal pair acceptability tests, which often fail to account for more complex patterns (cf. Gries 2003: Section 2.6.2 for discussion of such shortcomings in the area of syntax);
2. judgments are not gathered in an introspective way that relies on implicit knowledge and thus makes it difficult to validate and replicate findings;
3. instead, corpora
 - a. provide many instances rather than a few isolated judgments;
 - b. provide data from natural settings rather than ‘armchair’ judgments or responses that potentially reflect experimentally-induced biases;
 - c. provide co-occurrence data of many different kinds, i.e. not just those a particular researcher may consider important;
 - d. and thus, allow for bottom-up identification of relevant distinctions as well as for a more comprehensive description than is typically provided.

In this study, we will introduce a methodology that aims to provide the best of both worlds, i.e. a precise, quantitative corpus-based approach that yields cognitive-linguistically relevant results.

2. Methods

Our method is based on two key concepts. One is the notion of *ID tag* as proposed by Atkins (1987). The other is Hanks's (1996) notion of *Behavioral Profile*, which we extend from being restricted to complementation patterns and roles to include a comprehensive inventory of elements co-occurring with a word within the confines of a simple clause or sentence in actual speech and writing.

Our approach hinges on the assumption that the words or senses investigated are part of a network of words/senses. In this network, elements which are similar to each other are connected in such a way that the strength of the connection reflects the likelihood that the elements display similar behavior in other linguistic subdomains. The corpus-based method we will introduce focuses on co-occurrence information of symbolic units since (i) the symbolic unit is considered the basic unit within a cognitive linguistic approach and (ii) co-occurrences of this type are most easily accessible for a corpus-based approach.

The method involves the following four steps:

- i. the retrieval of (a representative random sample of) all instances of a word's lemma from a corpus;
- ii. a (so far largely) manual analysis of many properties of the word forms (i.e. the annotation of the ID tags);
- iii. the generation of a co-occurrence table;
- iv. the evaluation of the table by means of exploratory and other statistical techniques.

The first three of these steps are concerned with data processing, and will be dealt with in Section 2.1. The fourth step is concerned with how the resulting data can be evaluated meaningfully; it will be covered in detail in Section 2.2.

2.1 Data processing

Let us go over the data processing steps in somewhat more detail. The first step involves using a concordancing program, a programming language (e.g. R or Perl), or a corpus interface to retrieve (a subset of) all hits for the lemmata of a word or words of interest.²

In the second step, all hits are annotated for the ID tags one wishes to include in the analysis (cf. Section 4 below for discussion) in such a way that the results of the annotation process can be imported into spreadsheet software at a later stage. The range of ID tags that can be used is vast since virtually every linguistic level of analysis can be included. Table 1 provides a summary of ID tags that have been used so far.³

2. Note in passing that we use lemmata in order to be able to investigate whether particular inflectional forms behave differently from others. However, nothing in particular hinges on this decision and one might just as well base the study on the frequencies resulting from combining all inflectional forms of a lemma (cf. Gries to appear for discussion).

3. This list of ID tags results from our work on English and Russian. It is not exhaustive as far as senses are concerned and could be extended with additional ID tags (from the same domains or others such as phonology or pragmatics) or with ID tags manifested in other languages.

Table 1. Selective overview of (kinds of) ID tags and their levels

Kind of ID tag	ID tag	Levels of ID tag
morphological	tense	present, past, future
	mode	infinitive, indicative, subjunctive, imperative, participle, gerund
	aspect	imperfective vs. perfective
	voice	active vs. passive
	number	singular vs. plural
	transitivity	intransitive, monotransitive, copular, complex transitive
syntactic	sentence type	declarative, exclamative, imperative, interrogative
	clause type	main vs. dependent
	type of dependent clause	adverbial, appositive, relative, zero-relative, zero-subordinator, etc.
	semantic types of subjects, objects, etc.	concrete vs. abstract, animate (human, animal) vs. inanimate (event, phenomenon of nature, body part, organization/institution, speech/text) etc.
semantic	countability of nouns	count vs. mass
	properties of the process denoted by the verb	physical actions, physical perception, communication, intellectual activities, emotions, etc.
	controllability of actions	high vs. medium vs. no controllability
	adverbial/PP modification (if present)	temporal, locative, etc.
	negation	present vs. absent, attached to which element
lexical	collocates in precisely-defined syntactic slots (i.e. collexemes)	collocate ₁ , collocate ₂ , ..., collocate _n

Table 2. An excerpt from a co-occurrence table for *to run*

Citation	transitivity	morph. form	clause type	sense
Bert's now the priest who runs it	monotrans	present tense	depend	to manage
I will run out of money	intrans	infinitive	main	to lack
Troopers said the child ran into the path of a passing car	intrans	past tense	depend	to go very rapidly

The result of the second step is a table with co-occurrence information. In other words, each row contains one citation of the word in question, each column contains an ID tag and each cell contains the level of the ID tag for this citation. Table 2 contains an excerpt from the table used for the analysis of polysemous *run* in English (with examples from the ICE-GB). An analogous table for the investigation of near synonymous words would feature the near synonym in the last column (instead of the sense of a polysemous word).

In a third step, this table is prepared for quantitative analysis; this step consists of two phases. First, Table 2 is turned into a frequency table in a way that every row contains a level of an ID tag while every column contains a sense of the polysemous word

Table 3. Absolute co-occurrence frequencies of (levels of ID tags) and word senses

ID tag	level of ID tag	manage	lack	go very rapidly
transitivity	intransitive	0	12	191
	monotransitive	101	1	12
	copular	0	1	0
	complex transitive	0	0	0
morphological form	infinitive	25	1	43
	present tense	15	5	11
	present participle	23	4	54
	past tense	10	2	78
	past participle	28	2	11
	imperative	0	0	6

Table 4. Relative co-occurrence frequencies of (levels of ID tags) and word senses

ID tag	level of ID tag	manage	lack	go very rapidly
transitivity	intransitive	0	0.8571	0.9409
	monotransitive	1	0.0714	0.0591
	copular	0	0.0714	0
	complex transitive	0	0	0
morphological form	infinitive	0.2475	0.0714	0.2118
	present tense	0.1485	0.3571	0.0542
	present participle	0.2277	0.2857	0.2660
	past tense	0.0990	0.1429	0.3842
	past participle	0.2772	0.1429	0.0542
	imperative	0.0000	0	0.0296

or one word of the set of near synonyms; consequently each cell in the table provides the frequency of occurrence of the ID tags with the word/sense (cf. Table 3). The summed frequencies within each ID tag must be the same: for the sense *go very rapidly* this means that the sum of transitivity related ID tags (191+12) equals the sum of ID tags that capture morphological form (43+11+54+78+11+6).

In order to compare senses that occur at different frequencies, the absolute frequencies from Table 3 need to be turned into relative frequencies (i.e. within ID tag percentages; cf. Table 4).

In a quantitative, narrow sense of the term, Table 3 and Table 4 form the *behavioral profile* for a word/sense. In other words, each sense of a word or each near synonym within a semantic domain is characterized by one co-occurrence vector of within-ID tag relative frequencies.⁴ It is worth pointing out that this approach is compatible with at least two of the criteria of the principled-polysemy framework, namely the concept-elaboration criterion, positing distinct syntagmatic co-occurrence relations, and the grammatical criterion, positing distinct grammatical constructions. In fact, one could even say that our

4. Thus, the notion of behavioral profile is not related to the concept of profiling in cognitive grammar.

behavioral profile approach is based on taking these criteria and their manifestations as seriously as present-day corpora and efficiency demands allow. The following section will explain how behavioral profiles can be evaluated.

2.2 Evaluation

The vector-based behavioral profile can be subjected to a variety of quantitative approaches for further evaluation. There exist monofactorial and/or pairwise approaches as well as more comprehensive techniques that account for more complex multifactorial patterns. In Section 2.2.1, we will introduce some monofactorial methods, which will then be exemplified in more detail on the basis of the English verb *run* in Section 3.1. In Section 2.2.2, we will introduce a multifactorial cluster-analytic method, the application of that method to Russian verbs that express 'try' will be exemplified in Section 3.2.

2.2.1 *Monofactorial evaluation*

The most straightforward ways of analyzing behavioral profiles are looking at both token frequencies and type frequencies. Let us start with token frequencies. A useful first strategy is identifying in the corpus the most frequent senses of the word(s) one is investigating or the most frequent words within the semantic field studied. So far, our discussion has been non-committal with respect to the type of corpus investigated, but depending on the corpus the identification of the most frequently occurring word(s) or sense(s) may license different conclusions. In a general synchronic corpus, overall token frequency may be correlated with the degree of entrenchment of a word sense or of a word in a semantic field as well as its prototypicality (cf. Geeraerts 1988: 222; Winters 1990). In an acquisition corpus, tracking high percentages of senses and words across time and monitoring how they change over time may license conclusions about the ease of acquisition of senses and words as well as straightforward ways of semantic extension. In a diachronic corpus, the same procedure allows us to concentrate on the historical primacy of senses or words as well as on possible paths of extension and grammaticalization. While corpus-based work has been carried out in all of these areas, it typically takes a slightly more restricted stance in that the behavioral profiles entering into the analyses tend to be confined to many fewer ID tags than we propose.

While the inspection of frequencies is ultimately based on high *token* frequencies of particular ID tags, inspecting the *type* frequencies of ID tags is also revealing. Type frequencies should be 'normalized', i.e. the number of ID tags should be corrected against the overall frequency of occurrence of the sense or word (for instance, by dividing the number of observed ID tag types by the frequency of occurrence of that sense or word). The word senses or words with the highest number of non-zero values, i.e. the highest number of different ID tags, found in the behavioral profile correspond to unmarked senses or words since these senses/words exhibit the fewest restrictions concerning the range of ID tags applicable to them. Again, this may be an interesting finding in itself, as there is a positive though by no means absolute correlation between markedness and prototypicality (cf. Lakoff 1987: 60–61) which may be worth exploring. Yet, data of this type also allow the identification of exactly those cases where the co-occurrences of

senses/words and particular ID tags seem impossible, which in turn invites interesting semantic conclusions. Croft (1998: 169), for example, argues that disjoint syntactic-semantic distributions of otherwise similar senses support splitting senses as opposed to lumping them together.⁵

In addition, the distributional form in which the data come allows for more technical approaches from computational linguistics, where vectorized data underlie work on the semantic similarity of words, document clustering, and information retrieval (cf. Manning and Schütze 2000: Section 8.5). Moreover, the behavioral profile facilitates quantifying (and, thus, rank-ordering) senses or words in terms of their pairwise similarity (for more complex approaches, cf. Section 2.2.2 below); this goal can be achieved by computing any of several available similarity measures for vectors such as standard correlation coefficients, cosines, or other more complex indices. For example, network-inspired analyses of polysemous words require decisions as to where to locate senses in the network, and one way of approaching this issue is to first determine the highest pairwise similarities of the senses/words in question and then connect them to the senses/words they are most similar to.

One common characteristic of all of the above listed techniques is that they are monofactorial. That is, they are built either on vectors, i.e. one-dimensional distributions of percentages, or on pairwise similarities between vectors. However, the behavioral profile approach we are promoting here has more to offer and in the next section we will outline how multifactorial techniques can be brought to bear on the issues raised so far.

2.2.2 Multifactorial evaluation

There is quite a number of multifactorial techniques that could be applied to extract relevant information from behavioral profiles; we will restrict our attention to the exploratory technique of hierarchical agglomerative cluster analysis since it has been applied most frequently in related domains (cf. Manning and Schütze 2000: Chapter 14 for examples and discussion).⁶ The kind of cluster analysis that we advocate can be seen as consisting of three different steps, which we will discuss in turn.

The first step of the analysis consists of the hierarchical agglomerative cluster analysis proper of the joint behavioral profiles under investigation. Hierarchical agglomerative cluster (HAC) analysis is a family of methods that aims at identifying and representing (dis)similarity relations between different items; cf. Kaufman and Rousseeuw (1990)

5. It may likewise be possible to use the distributional data for exploring the acquisition of senses/words in a way complementing the approach mentioned above: equally frequent senses/words may differ in terms of their co-occurrence restrictions. A viable question would then be whether the more widely distributed senses/words give rise to extension of the category earlier than the more restricted ones. A similar logic applies to the case of diachronic corpora; cf. Bybee and Thompson (1997) for a pertinent discussion on type vs. token frequencies.

6. Techniques other than cluster analyses that can be applied to the kind of data discussed are singular value decomposition techniques (such as factor analyses or LSA), techniques for the multidimensional analysis of frequency tables (such as loglinear analysis or configural frequency analysis) and tree-based classification methods.

for a general discussion of clustering. Usually, clustering is performed on the basis of variables that characterize the items or on the basis of a (dis)similarity matrix of the items. In the TRY case, 1,585 corpus extractions that include examples for all nine near-synonymous verbs were tagged for 87 variables, i.e. our ID tags (a selection of which is shown in Table 1). Assigning ID tags to extractions resulted in the dataset represented in Table 4 above. Table 4 needs to be turned into a similarity/dissimilarity matrix, however, which can be done by means of a suitable similarity/dissimilarity measure. Since there are several measures available which differ along one or more parameters and thus may yield different cluster solutions, it is impossible to recommend any one specific measure: what is most suitable in one case (or with one set of assumptions one has about the data) may not work in another. It is probably fair to say that Euclidean distances (or squared Euclidean distances if one wants to ‘punish’ outliers) are among the most widely used measures in linguistic analyses.

Once the similarity/dissimilarity matrix has been generated, an amalgamation strategy has to be selected. An amalgamation strategy is an algorithm that defines how the elements that need to be clustered will be joined together on the basis of the variables or ID tags that they were inspected for. Again, the same caveats apply as for the generation of the similarity/dissimilarity matrix. One of the most widely used amalgamation strategies is Ward’s rule: it is conceptually similar to the logic underlying analysis of variance and typically yields moderately sized clusters.⁷

The result of such an analysis is a hierarchical tree diagram representing, in the ideal case, several relatively easily distinguishable clusters that are characterized by high within-cluster similarity and low between-cluster similarity. Often, the information gleaned from such a diagram is revealing in itself since the diagram summarizes conveniently what a human analyst could hardly discern given the complexity of a multifactorial data set.

The second step of the analysis consists of a detailed analysis of the clustering solution which (i) assesses the ‘cleanliness’ of the tree diagram and (ii) focuses on precisely those kinds of similarity that emerge most clearly from the tree diagram: between-cluster similarity and within-cluster similarity (cf. Backhaus et al. 2003: Chapter 8). As to the former, by a variant of the *F*-test also used in analyses of variance, it is possible to determine how homogenous the obtained clusters are. Obviously, the more homogenous the clusters are, the easier the interpretation of the between-cluster differences will be. As to the latter, it is possible to use *t*-values to determine which of the ID tags used reflect between-cluster differences best. More specifically, one can compute a *t*-value for each ID tag for each cluster such that a positive/negative *t*-value of an ID tag for a cluster indicates that this ID tag is respectively over-represented or under-represented in that cluster. This way, it is, for instance, possible to identify ID tags that have a positive *t*-value in one cluster and negative values in all other clusters, thus revealing the scales of variation that matter most for the clustering solution.

7. An alternative possibility is the choice of a phylogenetic clustering algorithm (cf. Felsenstein 2005 for an implementation), which does not require all elements that need to be clustered to be merged into a single root.

The third and final step consists of a similarly detailed analysis of the within-cluster differences. The fact that a cluster analysis has grouped together particular senses/words does not necessarily imply that these senses or words are identical or even highly similar – it only shows that these senses/words are more similar to each other than they are to the rest of the senses/words investigated. By means of standardized *z*-scores, one can tease apart the difference between otherwise highly similar senses/words and shed light on what the internal structure of a cluster looks like.

While the discussion has been relatively abstract so far, we will now present several examples to illustrate how the methods introduced above can be put to use.

3. Examples

In this section, we will discuss examples from a case study on an extremely polysemous English verb (Section 3.1) and from a case study of nine near synonymous Russian verbs (Section 3.2).

3.1 Polysemy: The English verb *run*

The examples to be discussed in this section are taken from Gries (2006) that deals with the highly polysemous English verb *run*.⁸ The analysis is carried out using 815 citations of the verb lemma *run* from two corpora; each citation was coded for the senses they instantiate within their respective contexts as well as for 252 ID tags of the types given in Table 1; many of the ID tags in this study code the presence/absence of particular collocates.

Let us begin with the issue of how one-dimensional vectors (frequency distributions) can be exploited to address the question of prototypical word senses, an issue where corpus data can be applied in a versatile way. In this case, the corpus data clearly single out one sense, namely the sense ‘fast pedestrian motion’. This is the sense that is

- diachronically primary: together with ‘*flow*’ it is the earliest attested sense;
- diachronically primary for the zero-derived noun *run*;
- synchronically most frequent in the analyzed corpora;
- synchronically most frequent for the zero-derived noun in the analyzed corpora;
- acquisitionally primary in the sense of being acquired earliest;
- acquisitionally most frequent (counts from data for Abe, Adam, Eve, Naomi, Nina, Peter, and Sarah from the CHILDES database; cf. MacWhinney 2000);
- combinatorially least constrained in the analyzed corpora (given its number of ID tags normalized against frequency of occurrence).

8. Cf. Langacker (1988) and Taylor (1996, 2000) for cognitive-linguistic but methodologically very different studies of the verb *run*.

Vectors can likewise be used to identify disjoint distributions, as the examples of ‘fast pedestrian motion’ and ‘escape’ show. Applying Croft’s (1998) logic, for example, one would not consider the senses instantiated in (1a) and (1b) as different merely because their PPs highlight different landmarks. This is so because there are also examples like (2) in which the two kinds of PPs – SOURCE and GOAL – co-occur, showing that the distribution of the PPs is not disjoint.

- (1) a. and we ran back [_{GOAL} to my car]
 b. Durkin and Calhoun came running [_{SOURCE} from the post]
- (2) I ran [_{SOURCE} from the Archive studio] [_{GOAL} to the Start The Week studio]

However, there are other senses, intuitively very similar, which are likely candidates for being lumped together. For example, there are two senses that could both be paraphrased as ‘escape’, but one of them involves moving away from something undesirable while the other involves moving away to engage in a romantic relationship. Interestingly, the former (see (3)) is attested with a SOURCE but not with a comitative argument whereas the latter (see (4)) is attested with a comitative but not with a SOURCE although both unattested combinations are conceivable.

- (3) He wanted to know if my father had beaten me or my mother had run away
 [_{SOURCE} from home]
- (4) If Adelia had felt about someone as H. felt about C., would she have run away
 [_{COMITATIVE} with him]?

While the results of a corpus-based application of the criterion of disjoint distribution are certainly dependent on sample sizes, they indicate – in the absence of evidence to the contrary – that the two ‘escape’ senses should not be lumped together. Once it has been decided to keep these senses separate the question arises of where to connect them to the rest of the network. One possible point of connection would be the sense of ‘fast pedestrian motion’. Yet, not all the instances of the ‘escape’ senses imply fast pedestrian motion: some merely imply ‘fast motion’ or only ‘motion’. ‘Motion’ would therefore also be a plausible candidate sense for the connection. This issue can be solved by making use of the information contained in the behavioral profile for each sense. Pearson product moment correlations were computed for all pairs of senses in order to determine the average correlation of all senses but also to find out which of the three candidate senses are most similar to the two ‘escape’ senses that need to be connected. While the overall average correlation (after Fisher Z transformation) was moderate ($r = 0.545$), the average correlation of the two ‘escape’ senses and the three ‘motion’ senses was considerably higher ($r = 0.848$), supporting the intuition that these senses are in fact closely related, at least much more than they are related to the multitude of other senses that *run* can have. When the question of where to attach the two ‘escape’ senses was investigated using a smaller set of ID tags (omitting collocation-based ID tags lest individual collocates distort the picture), a surprisingly clear answer emerged. The two ‘escape’ senses were significantly more similar to ‘fast pedestrian motion’ than to the other two senses, which in turn did not differ significantly from each

other. This result provides evidence for attaching the two 'escape' senses to the prototypical sense as opposed to the two slightly more general senses.⁹

So far the examples presented involved only monofactorial data (for considerations of space, the cluster-analytic results presented in Gries 2006 are not discussed here). The following section will provide detailed exemplification of how cluster analyses and their follow-up investigation can be useful for the lexical semanticist.

3.2 Near synonymy: Russian verbs meaning *try*

In this section, based on Divjak and Gries (2006), we show how clustering behavioral profiles and evaluating clusters and verbs in terms of *t*-values and *z*-scores provide us with scales of variation for describing and distinguishing near synonyms in a fine-grained lexical semantic analysis. Divjak and Gries (2006) analyze 1,585 sentences each containing one out of nine Russian verbs that, in combination with an infinitive, express *try*. Since the verbs in question differ strongly in terms of their frequencies, the sentences were culled from several sources, keeping the genre constant: the Amsterdam corpus, the Russian National Corpus, and the WWW (cf. Divjak and Gries 2006: 54, note 6 for detailed discussion of the sampling procedure); Table 5 sketches the composition of the data set.

All 1,585 sentences were annotated for 87 ID tags; as a result, for each of the nine verbs a behavioral profile vector was obtained of the sort exemplified in Table 4. This dataset was analyzed using a hierarchical agglomerative cluster analysis (similarity metric: Canberra; amalgamation strategy: Ward), resulting in the dendrogram presented in Figure 1. The tree plot shows what is similar and what is different: items that are clustered or amalgamated early are similar, and items that are amalgamated late are rather dissimilar.

For example, it is obvious that *pytat'sja* and *starat'sja* are much more similar to each other than, say, *probovat'* and *norovit'*, which are only linked in the last overarching cluster. At the same time, the plot gives an indication of how independent the clusters are: the larger the distance between different points of amalgamation, the more autonomous the earlier verb/cluster is from the verb/cluster with which it is merged later. In the present case, the plot clearly consists of three clusters.

Table 5. Composition of the dataset analyzed in Divjak and Gries (2006)

Verb	N (AC/RNC/Web)	Verb	N (AC/RNC/Web)
<i>probovat'</i>	246 / - / -	<i>poryvat'sja</i>	31 / 88 / -
<i>pytat'sja</i>	247 / - / -	<i>tščit'sja</i>	21 / 30 / 21
<i>starat'sja</i>	248 / - / -	<i>pyžit'sja</i>	- / - / 98
<i>silit'sja</i>	57 / 185 / -	<i>tužit'sja</i>	- / - / 53
<i>norovit'</i>	112 / 148 / -		

9. Of course, this method is not restricted to cases where one sense needs to be attached to only one other sense. In cases where multiple attachments are desired, the correlations can still be used to rank or delimit the candidate set of senses to which another sense can be reasonably attached. Also, nothing hinges on the choice of the Pearson product moment correlation: as indicated above, other measures could be employed; in this particular case, the cosine measure was also tested and yielded the same conclusions.




Figure 1. Dendrogram for tentative verbs in Russian

A cognitive approach to language and particularly to the incorporation of knowledge about human categorization mechanisms into linguistics provides interesting perspectives for a unified interpretation of the data. On the cognitive linguistic approach, (linguistic) categories may exhibit prototype effects and instantiate radial networks of related expressions with semantically motivated connections (Lakoff 1987: Chapter 6).¹⁰ In order to investigate the nature of the three categories suggested by the dendrogram more thoroughly, between- and within-cluster similarities and differences were inspected using *t*-values and *z*-scores (cf. above); limitations of space permit only a selection of the results to be discussed.

The first cluster groups together [*pytat'sja* and *starat'sja*] and *probovat'*. All verbs in this cluster are more easily used in the main clause ($t = 0.821$) than verbs from the other two clusters. Although all three verbs exist in the imperfective and perfective aspect and do occur in both aspects, variables that include reference to the perfective aspect (i.e. refer to past and future events) are three times more frequent in the top 25 *t*-scores that are positive for this cluster and negative for other clusters (*t*-values range from 0.667 to 1.201). In addition, the infinitive that follows the tentative verb is more often negated ($t = 0.702$) and expresses physical activities ($t = 0.599$), events that are figurative extensions of motion events ($t = 0.465$) or involve setting a theme/patient into motion ($t = 0.4$). Finally, strongly

10. Although the HAC dendrogram presented in Figure 1 can be manually transformed into a radial network representation, Divjak and Gries (2006) backed up their results by analyzing the distance matrix resulting from the behavioral profiles using a phylogenetic clustering algorithm, the Fitch program from the PHYLIP package (Felsenstein 2005). The results were for all practical purposes identical; cf. Divjak and Gries (2006: Section 3) for discussion.

attracted optional collocates express that the subject got permission to carry out the infinitive action (using *pust'*, $t = 1.008$), that the attempt was untimely brought to a halt (with *bylo*, $t = 0.982$), that the subject was exhorted to undertake an attempt ($t = 0.832$) and that the intensity with which the attempt was carried out was reduced ($t = 0.667$).

In the middle, there is a cluster that unites the imperfective verbs [*tščit'sja* and *pyžit'sja* and *tužit'sja*]. All three verbs lack a perfective counterpart and prefer the present tense more than verbs in the two other clusters ($t = 1.047$ for present tense with a perfective infinitive and $t = 0.711$ for the present tense followed by an imperfective infinitive). Among the most strongly represented variables we encounter the verbs' compatibility with inanimate subjects, both concrete and abstract (t ranges from 1.108 to 1.276), as well as with groups or institutions ($t = 1.297$). Actions expressed by the infinitive are physical ($t = 0.176$), affect a theme/patient ($t = 0.352$), are metaphorical extensions of physical actions ($t = 0.999$), or physical actions affecting a theme/patient ($t = 0.175$). Focus is on the vainness ($t = 0.962$ for vainness combined with intensity) of the durative effort ($t = 0.750$ for duration adverbs).

The third cluster, amalgamated last into the overarching cluster, consists of [*norovit'* and *poryvat'sja*] and *silit'sja*. These verbs prefer to occur as participles (t 's range from 0.632 to 1.214). The infinitive actions that are attempted express a type of physical motion ($t = 0.924$) that is often not controllable ($t = 0.548$). The action can be carried out by an inanimate subject ($t = 0.809$ for phenomena of nature and $t = 0.774$ for bodyparts) and are often repeated (t ranges from 0.678 to 1.092). If the attempt remains unsuccessful, both external ($t = 0.627$) and internal ($t = 0.429$) reasons are given for the failure.

Apart from between-cluster differences that are revealed by means of t -scores, z -values make within-cluster similarities and differences visible. As an illustration, let us look at the three most frequently used verbs, i.e. the verbs in the first cluster [*pytat'sja* and *starat'sja*] and *probovat'*. The two verbs that are clustered first, *pytat'sja* and *starat'sja*, resemble each other to a large extent, yet a close inspection of their distributional properties reveals that *pytat'sja* is more strongly attracted to occurring in the past tense (with z 's ranging from 1.092 to 1.155, all with perfective infinitives) whereas *starat'sja* is relatively more often found in the present tense ($z = 1.153$ with imperfective infinitives). *Pytat'sja* is not particularly attracted to weakly controllable actions ($z = -1.097$) whereas *starat'sja* avoids controllable actions ($z = -1.049$). *Starat'sja* combines, among other things, with passive perception verbs ($z = 1.134$), whereas *pytat'sja* goes well with mental activities ($z = 1.139$). *Starat'sja* is frequently found with a negated infinitive ($z = 1.151$), thus indicating that the subject is avoiding an event that might take place. Easiest to interpret is the verbs' preference for different adverbs: *starat'sja* is most strongly characterized by adverbs that express repetitive duration (*vsě vremja*, $z = 1.155$), reduced intensity ($z = 1.155$), and intensity ($z = 1.101$), whereas *pytat'sja* prefers repetition ($z = 1.111$). In other words, if one has already applied *pytat'sja* without success, a possible way to achieve the desired result despite the initial failure is by using what is encoded in *starat'sja* (cf. (5)).

- (5) Он убрал Мазера и Леоновича, постарается то же проделать с Казаковым (уже пытался), и весьма возможно, с Соя-Серко.
[Ф. Незнанский, Ярмарка в Сокольниках]
'He took away Mazer and Leonovič, is trying (hard) [starat'sja] to do the same with Kazakov (he has already tried [pytat'sja]), and it is very likely, with Soja-Serko.'

Added to [pytat'sja and starat'sja] is the verb *probovat'* that is rather dissimilar. This verb occurs preferably in a main clause ($z = 1.127$), and is not typically found in declarative clauses ($z = -1.148$). Tags that refer to perfective aspect receive the highest z -scores for [probovat'], ranging from 1.003 to 1.155. Although all three verbs in this cluster have a perfective counterpart formed by means of the delimitative prefix *po-*, *po/probovat'* significantly prefers the perfective aspect in 74.8% of all examples while *pytat'sja* and *starat'sja*, by contrast, significantly prefer the imperfective aspect, i.e. in 79.6% and 83% of all cases respectively ($\chi^2 = 222.72$; $df = 2$; $p < 0.001$, Cramer's $V = 0.548$). Related to the more frequent use of perfective forms is the possibility of locating the attempt in the future ($z = 1.003$ for combinations with imperfective infinitives and $z = 1.044$ with perfective infinitives), as well as a considerable relative dispreference for the present tense (z 's ranges from -0.632 to -1.154). Finally, *probovat'* is the only verb that is often found in the imperative mode (with z 's ranging from 1.092 to 1.134). In interpretive terms, the node [probovat'] uses the perfective to present each try as a completed entity. This allows the subject to change method or strategy between attempts, which might be what makes this verb resemble experiments (cf. Wierzbicka 1988: 309; Apresjan et al. 1999: 304). An experimental attempt is also demanded more easily from another person than attempts that require long and/or intense effort, hence the higher frequency of the imperative and attraction of exhortative particles ($z = 1.121$). Failure can be attributed to internal and external factors alike (4.9%, $z = 1.155$ and 11%, $z = 1.151$). In all, *probovat'* seems to be less intensive than *pytat'sja* (and *starat'sja*), as example (6) shows.

- (6) Бим уже пробовал на нее наступить, но пока еще так, немножко – только пробовал. [Г. Трупольский. Белый Бим черное ухо]
'Bim had already tried [probovat'] to step on her, but just like that, a little bit, he had only tried [probovat'].'

The multifactorial evaluation we propose comprises a set of both exploratory and hypothesis-testing statistical techniques for analyzing corpus-based behavioral profiles. We have illustrated how, on the basis of these results, the internal structure of a cluster of near synonymous verbs can be laid bare and the verbs in those clusters can be compared.

4. Conclusion

We hope to have shown that behavioral profiles and the proposed methods for their evaluation are valuable for the analysis of polysemous and near synonymous items in particular as well as for lexical-semantic research in general. Moreover, behavioral profiles provide an ideal starting point for research concerning interfaces between different levels of linguistic analysis, e.g. the syntax-lexis interface, and offer a wealth of usage-based evidence

for cognitive linguistic theorizing concerning network representations, prototypicality of senses, sense-distinctions and the polysemy-homonymy discussion to name but a few. In addition, results of this type may also be relevant for researchers from neighboring disciplines, such as psycholinguistics: behavioral profiles can be used in formulating and evaluating hypotheses concerning the interaction between grammar and lexicon in language acquisition as well as with respect to the mental reality of radial categories (cf. Divjak and Gries 2008). Conveniently, a program for converting annotated data into behavioral-profile vectors and computing cluster-analytic statistics is now available (cf. Gries 2008).

Our plea for a corpus-based approach does not imply adherence to a fully automated approach, however. At present there is no reliable way for assigning (many) ID tags automatically and neither can a machine interpret statistical results. Although human intervention rules out complete objectivity, we do claim that our methodology is more objective than many others currently available. The proposed approach requires all information entering into the analysis to be made explicit: it is necessary to define and operationalize every ID tag since it is only through frequency counts of ID tags that information can be included. In other words, our method helps to minimize the share of subjective, implicit knowledge. In addition, while the choice of ID tags to be included in the analysis and the subsequent interpretation of the results contain elements of subjectivity – as does, if to a lesser degree, the annotation/coding of the dataset – a substantial part of the analysis is entirely objective. For example, an analyst cannot simply select parameters or ID tags for interpretation *ad libitum*, but is strongly constrained by the statistical results which were arrived at in an objective and replicable way (a hierarchical agglomerative cluster analysis can be defined precisely in terms of its mathematical settings). Thus, if, say, a *t*-score does not differentiate (significantly) between clusters, the analyst cannot belabor its importance however much his theoretical commitment would require him to. For these reasons we submit that the behavioral profile approach as outlined above is an improvement over many other methodological tools in the domain of lexical semantics in general and cognitive lexical semantics in particular.

References

- Apresjan, J. D., Boguslavskaja, O. J., Levontina, I. V., Uryson, E. V., Glovinskaja, M. J. & Krylova, T. V. 1999. *Новый объяснительный словарь синонимов русского языка. Volume I*. Moskva: Škola “Jazyki Russkoj Kul'tury”.
- Atkins, B. T. S. 1987. Semantic ID tags: Corpus evidence for dictionary senses. *Proceedings of the Third Annual Conference of the UW Centre for the New Oxford English Dictionary*: 17–36.
- Backhaus, K., Erichson, B., Plinke, W. & Weiber, R. 2003. *Multivariate Analysemethoden: Eine Anwendungsorientierte Einführung*. Berlin: Springer.
- Bybee, J. & Thompson, S. 1997. Three frequency effects in syntax. *Proceedings of the Twenty-third Annual Meeting of the Berkeley Linguistics Society*: 378–388.
- Corston-Oliver, M. 2001. Central meanings of polysemous prepositions: Challenging the assumptions. Poster presented at the ICLC 2001, University of California at Santa Barbara.
- Croft, W. 1998. Linguistic evidence and mental representations. *Cognitive Linguistics* 9 (2): 151–173.
- Croft, W. 2001. *Radical Construction Grammar*. Oxford: Oxford University Press.
- Cruise, D. A. 1986. *Lexical Semantics*. Cambridge: Cambridge University Press.

- Divjak, D. 2004. Degrees of Verb Integration. Conceptualizing and Categorizing Events in Russian. PhD Dissertation, K.U. Leuven (Belgium).
- Divjak, D. 2006. Ways of intending: Delineating and structuring near synonyms. In *Corpora in Cognitive Linguistics: Corpus-Based Approaches to Syntax and Lexis*, St. Th. Gries & A. Stefanowitsch (eds), 19–56. Berlin: Mouton de Gruyter.
- Divjak, D. & Gries, St. Th. 2006. Ways of trying in Russian: Clustering behavioral profiles. *Corpus Linguistics and Linguistic Theory* 2 (1): 23–60.
- Divjak, D. & Gries, St. Th. 2008. Clusters in the mind? Converging evidence from near synonymy in Russian. *The Mental Lexicon* 3 (2): 188–212.
- Evans, V. 2005. The meaning of *time*: Polysemy, the lexicon and conceptual structure. *Journal of Linguistics* 41 (1): 33–75.
- Felsenstein, J. 2005. *Phylogeny Inference Package PHYLIP 3.65*, <http://www.phylip.com>.
- Geeraerts, D. 1985. Preponderantieverschillen bij bijna-synoniemen. *De Nieuwe Taalgids* 78: 18–27.
- Geeraerts, D. 1988. Where does prototypicality come from? In *Topics in Cognitive Linguistics*, B. Rudzka-Ostyn (ed.), 207–229. Amsterdam: John Benjamins.
- Gibbs, R. W. & Matlock, T. 2001. Psycholinguistic perspectives on polysemy. In *Polysemy in Cognitive Linguistics*, H. Cuyckens & B. Zawada (eds), 213–239. Amsterdam: John Benjamins.
- Gries, St. Th. 2003. *Multifactorial Analysis in Corpus Linguistics: The Case of Particle Placement*. London: Continuum.
- Gries, St. Th. 2006. Corpus-based methods and cognitive semantics: The many meanings of *to run*. In *Corpora in Cognitive Linguistics: Corpus-Based Approaches to Syntax and Lexis*, St. Th. Gries & A. Stefanowitsch (eds), 57–99. Berlin: Mouton de Gruyter.
- Gries, Stefan Th. 2008. BP. A program for R (for Windows).
- Gries, St. Th. To Appear. Corpus data in usage-based linguistics: What's the right degree of granularity for the analysis of argument structure constructions? In *Expanding Cognitive Linguistic Horizons*, M. Brdar, M. Ž. Fuchs & St. Th. Gries (eds).
- Hanks, P. 1996. Contextual dependency and lexical sets. *International Journal of Corpus Linguistics* 1 (1): 75–98.
- Kaufman, L. & Rousseeuw, P. J. 1990. *Finding Groups in Data*. New York: John Wiley.
- Kishner, J. M. & Gibbs, R. W. 1996. How *just* gets its meanings: Polysemy and context in psychological semantics. *Language and Speech* 39 (1): 19–36.
- Klein, D. E. & Murphy, G. L. 2001. The representation of polysemous words. *Journal of Memory and Language* 45 (2): 259–282.
- Klein, D. E. & Murphy, G. L. 2002. Paper has been my ruin: Conceptual relations of polysemous senses. *Journal of Memory and Language* 47 (4): 548–570.
- Kreitzer, A. 1997. Multiple levels of schematization: a study in the conceptualization of space. *Cognitive Linguistics* 8 (4): 291–325.
- Lakoff, G. 1987. *Women, Fire, and Dangerous Things*. Chicago: University of Chicago Press.
- Langacker, R. W. 1988: A usage-based model. In *Topics in Cognitive Linguistics*, B. Rudzka-Ostyn (ed.), 127–161. Amsterdam: John Benjamins.
- MacWhinney, B. 2000. *The CHILDES Project: Tools for Analyzing Talk*. Mahwah, NJ: Lawrence Erlbaum.
- Manning, C. D. & Schütze, H. 2000. *Foundations of Statistical Natural Language Processing*. Cambridge, MA: M.I.T. Press.
- Mondry, H. & Taylor, J. R. 1992. On lying in Russian. *Language and Communication* 12 (2): 133–143.
- Norvig, P. & Lakoff, G. 1987. Taking: a study in lexical network theory. In *Proceedings of the Thirteenth Annual Meeting of the Berkeley Linguistics Society*: 195–206.
- Rice, S. 1996. Prepositional prototypes. In *The Construal of Space in Language and Thought*, M. Pütz & R. Dirven (eds), 135–165. Berlin: Mouton de Gruyter.
- Sandra, D. & Rice, S. 1995. Network analyses of prepositional meaning: Mirroring whose mind – the linguist's or the language user's? *Cognitive Linguistics* 6 (1): 89–130.

- Stefanowitsch, A. & Gries, St. Th. 2003. Collostructions: Investigating the interaction between words and constructions. *International Journal of Corpus Linguistics* 8 (2): 209–243.
- Taylor, J. R. 1996. On running and jogging. *Cognitive Linguistics* 7 (1): 21–34.
- Taylor, J. R. 2000. Approaches to word meaning: the network model (Langacker) and the two-level model (Bierwisch). In *The Lexicon-Encyclopedia Interface*, B. Peeters (ed.), 113–141. Amsterdam: Elsevier.
- Taylor, J. R. 2003. Near synonyms as co-extensive categories: *high* and *tall* revisited. *Language Sciences* 25 (3): 263–284.
- Tyler, A. & Evans, V. 2001. Reconsidering prepositional polysemy networks: The case of *over*. *Language* 77 (4): 724–765.
- Wierzbicka, A. 1988. *The Semantics of Grammar*. Amsterdam: John Benjamins.
- Winters, M. 1990. Toward a theory of syntactic prototypes. In *Meanings and Prototypes: Studies in Linguistic Categorization*, S. L. Tsohatzidis. (ed.), 285–306. London: Routledge.

Polysemy, syntax, and variation

A usage-based method for Cognitive Semantics

Dylan Glynn

1. Introduction

The study of polysemy has a venerable tradition in Cognitive Linguistics.¹ Since the pioneering of work of Dirven (1981), Radden (1981), Brugman (1983), Lindner (1983), and Vandeloise (1984), the Lexical Network approach to spatial prepositions has proven to be one of the most important contributions of the cognitive paradigm.² However, at both a theoretical and methodological level, this network approach to sense variation has recently come under fire. This study examines a methodological proposal that answers some of the criticisms that a cognitive approach to polysemy faces. The basic premise is to conserve the network model, but to complement this with another method: a corpus-driven quantified and multifactorial method. Such an approach employs a kind of componential analysis that identifies clusters of features across large numbers of speech events. In other words, rather than analyse the possible meanings of a lexeme, a polysemic network should “fall out” from an analysis that identifies clusters of the cognitive-functional features of a lexeme’s usage. These features do not in any way resemble those of the Structuralist componential analyses, since they are not based on a hypothetical semantic system, but describe instances of real language usage and are based upon encyclopaedic semantics of that language use in context.

This usage-based approach is gaining wide currency in Cognitive Linguistics, attested by the wide range of edited volumes dedicated to the subject (Gries and Stefanowitsch 2006; Stefanowitsch and Gries 2006; Zeschel 2008; Glynn and Fischer in press; Newman and Rice in press; Glynn and Robinson forthcoming). We will refer to this approach as the Quantitative Multifactorial method. The discussion begins by briefly covering the current state of the art of polysemy study in cognitive linguistics, where we underline the need for the implementation of this method and how it may cooperate with existing analytical models. In Section 2, the discussion moves to the description of the Quantitative Multifactorial method, and finally the third section examines a case study of the English lexeme *hassle*.

-
1. Thanks are due to Tine Breban, Sofie van Gijssel, and Koen Pleveots. All shortcomings are my own.
 2. Some variations upon and applications of this approach include Schulze (1988, 1994), Hottenroth (1991), Geeraerts (1992), Cuyckens (1993, 1994, 1995), Boers (1996), Bellavia (1996), and Meex (2001).

2. Semasiology: Vagueness, polysemy, and the lexical network

Two groundbreaking studies demonstrate inherent weaknesses in what had become the standard radial model or lexical network approach to the identification sense variation in Cognitive Linguistics. These studies, Sandra and Rice (1995) and Tyler and Evans (2001), have led to a fundamental reconsideration of the conceptual reality represented by network modelling, bringing the validity of such an approach into question. The first study, by Sandra and Rice (1995), led to a debate that questions the psychological reality of sense distinctions proposed by the method and, through the use of psycholinguistic experimentation, raises serious doubts about the validity of the results.³ The second, more theoretical study by Tyler and Evans (2001) uses the very logic that led to the cognitive modelling of sense networks to demonstrate that the network model is flawed.

Thus far, solutions designed to resolve the shortcomings of the lexical network model decline into three approaches.⁴ Firstly, following the psycholinguistic tests of Sandra and Rice (1995), one may turn to an experimental solution to the unverifiable nature of intuition in sense identification and distinction. To these ends, attempts at developing psycholinguistic methods of testing have been pursued. This approach, examples of which include Cuyckens et al. (1997), Sandra and Cuyckens (1999), and Rice et al. (1999), makes the assumption that psycholinguistic experiments are a more reliable method than intuition for the identification of polysemic structures.

The second and third proposed solutions are analytical rather than methodological. One of these analytical solutions is to posit different “types” of polysemy. For instance, Kleiber (1983, 1999), Herskovits (1988), Vandeloise (1990), Deane (1988, 2006), and Glynn (2003, 2006b, 2006c) stress the need to distinguish functional and/or grammatical polysemy from conceptual polysemy.⁵ The assumption here is that, for example, a meaning extension may be the result of the interaction between the conceptual meaning associated with a lexeme and the meaning associated with a given grammatical category. Or, similarly, a socially determined function, such as implicature, may be the motivation behind a semantic extension.

The other analytical solution, developed by Tyler and Evans (2001, 2003), Evans and Tyler (2004a, 2004b), and Evans (2004, 2005), is named the Principled Polysemy Model and uses predetermined criteria to constrain the proposal of novel senses. This method improves on previous models precisely because sense distinction criteria are overtly specified. By doing this, the approach takes the bull by the horns and begins with the fundamental

3. The debate, exemplified by Croft (1998), Sandra (1998), and Tuggy (1999) follows from an earlier discussion over vagueness and polysemy, cf. Geeraerts (1993), and Tuggy (1993).

4. Other than the three approaches discussed here, there exist a range of other descriptive models more or less built on or derived from the lexical network model. Taylor (2003b) offers a summary of many of these. The crux of the issue is categorisation. Cf. Geeraerts (1989, 1990, 1997, 2000, 2006b), Lehrer (1990), Taylor (1993, 2003a), Cruse (1995), Lewandowska (2002, 2007), and Zlatev (2003).

5. This idea flows from the current trends in Structuralist linguistics that maintain a distinction between *langue* and *parole*. Cf. Fuchs (1987, 1991), Picoche (1994), Victorri and Fuchs (1996), and Rémi-Giraud and Panier (2003) for examples of contemporary Structuralist approaches to polysemy.

question: upon what grounds do we distinguish senses? However, both these analytical solutions (the principled and functional-conceptual) to solving the quandary of sense distinction suffer from the problems that the original Structuralist approach faced: intuition-based methods of investigation necessarily use intensional hermeneutic means for applying their criteria. For example, no matter how clearly set out, the three distinct sense criteria of the Principled Polysemy approach must be applied to language examples and human judgments must be made as to whether these criteria are met.

Nevertheless, every methodology has its strengths and weaknesses. The recent work using the Principled Polysemy approach has met with success because it focuses on precisely the problem at hand rendering the method's inherent weakness open to scrutiny and thus verification. It is thus that the Principled Polysemy approach seeks to resolve the tricky balance between what semantic variation is inherently associated with a given form and what semantic extensions this form–meaning pair may allow in given contexts. This, of course, brings us to the well-known vagueness versus polysemy debate. Principled Polysemy Modelling of Lexical Networks seeks precisely to elucidate the vague polysemic distinction through the proposal of criteria to distinguish “semantic elaborations” from “sanctioning senses” (Evans 2005: 38–40, 41–45).

A priori, all cognitive–functional research accepts axiomatically that “[word] meaning is highly context-sensitive, and thus mutable” (Evans 2005: 71). The question for a linguist faced with such an unstable object of study is not how to render sense variation more stable, but how to reveal structure in its variation. Zelinsky-Wibbelt (2000) poses this fundamental question for the study of polysemy: “Is polysemy a case of lexical representation or rather ... a case of contextual differentiation?” (Zelinsky-Wibbelt 2000: 144). Her discussion is surely amongst the most level-headed on the subject of how to distinguish entrenched sense variation from context-dependent variation. Theoretically in line with Tyler and Evans's (2001) premise, Zelinsky-Wibbelt phrases the problem as a methodological question: in the description of polysemy, “what should be represented at the level of the lexicon and what should be computed by contextual functions?” (Zelinsky-Wibbelt 2000: 145).

At play here is the role of real-time processing versus learnt-automated structure. Necessarily, the former is how one deals with context-dependent meaning production and the latter with entrenched meaning structure. Just as Tyler and Evans (2001: 726) cite universal cognitive processes as a means for explaining the interaction between entrenched lexical reference and contextual information, Zelinsky-Wibbelt poses her methodological question in the context of the basic cognitive process of construal. Here, she echoes the proposal of Kreitzer (1997) that context-dependent construal is an effective means for explaining much sense variation.⁶ She stresses that speakers “negotiate the reorganization of the same concept in potentially infinitely many ways. Vagueness ... represents the speaker's underlying continuum of knowledge from which their communicative ability of negotiating the relevant boundaries of meaning proceeds” (Zelinsky-Wibbelt 2000: 146). Tyler and Evans (2001: 726) stress this same point, arguing that the standard Lexical Network

6. Obviously, the role of different types of cognition in polysemy is a common theme in the literature. Within Cognitive Linguistics, Deane (1988) was probably the first to stress its importance.

approach to polysemy “fail[s] to distinguish between what is coded by a lexical expression and the information that must be derived from context. [Such approaches] fail to take account of meaning construction as a process.” They also follow Kreitzer and Zelinsky-Wibbelt in the use of construal to explain context-dependent polysemy. This brings us to the obvious conclusion that we should base the description of polysemy, and the structures that organise context-dependent construal of sense, on examples of usage, i.e. on examples of how speakers negotiate this basic cognitive and communicative phenomenon.

A Quantitative Multifactorial method aims to fulfil this role, not by adopting an analytical solution, but through a usage-based approach. In this, instead of identifying different senses, one looks for patterns of usage in terms of relative frequency. Its results may complement psycholinguistic testing to corroborate the hypothetical models of polysemy based on intuition. However, it is important to note here that corpus-driven and psychologically tested results cannot disprove intuitive results. This is because no corpus is large enough to account for every possibility and no set of informants large enough to represent the collective speakers of a speech community. Indeed, as we will see, intuitive methods, such as the Principled Polysemy Model, remain essential. However, usage-based quantitative results can complement the intuition-based study of polysemy by verifying core senses and offering information as to the relative semasiological structure of a lexeme. Since a corpus-driven method is inherently restricted to core senses, it cannot serve as a means for constraining the lexical network model. Nevertheless, depending on the representativity of the corpus, results can be argued to describe the most conceptually salient usages of a lexeme or its prototype structure, as well as capture differences between register and dialect. Moreover, and perhaps most importantly, it may reveal how such factors affect each other in the semantic structure of a lexeme. In other words, for a given register, speech situation, or dialect, one reading of a word may be more salient than another. This last point is one of the main advantages of a quantitative method. The multifactorial nature of language use may be rigorously described where, in intuition-based study, one cannot adequately account for the interaction of the different contextual parameters that affect meaning and usage.

Cognitive Linguistics is a usage-based approach to language (Langacker 1988, 2000) and, as such, must necessarily account for the complexities of language as a social phenomenon. This theoretical tenet means that a cognitive approach must necessarily consider extralinguistic parameters. Geeraerts’ plenary lecture at the 8th International Cognitive Linguistics Conference (published Geeraerts 2005) stressed the inevitability of Social Cognitive Linguistics, a point re-iterated by Croft (this volume) in the opening plenary of the first UK Cognitive Linguistics Conference.⁷ The question is not *if* we need to account for variables such as dialect, sociolect, and register, but *how* our analytical apparatuses can account for this complexity. The advantage of the Quantitative Multifactorial method is that this information is inherent to the analysis.

7. Kemmer’s keynote at the *First International Conference of the Swedish Cognitive Linguistics Association* joins the chorus arguing this point. Some other recent publications to press this argument include Tummers et al. (2005), Geeraerts (2006), Grondelaers et al. (2007), Heylen et al. (2008), and Glynn (in press).

The Quantitative Multifactorial method has, in fact, a well-established tradition in Cognitive Linguistics, but principally in the study of onomasiology and parasynonymy. Indeed, some of the earliest cognitive lexical research was quantitative, multifactorial, and usage-based. The early corpus-driven work of Dirven et al. (1982) and the elicitation-based research of Lehrer (1982) are excellent examples of this method. The approach slowly gained momentum with studies such as those of Zelinsky-Wibbelt (1986, 1993), Rudzka-Ostyn (1989, 1995), Schmid (1993), Geeraerts et al. (1994), Atkins (1994), Lemmens (1998), and Geeraerts (1999).

It is in light of the success of this research that Cognitive Linguistics has recently seen a blossoming in quantitative corpus-driven methodology. The use of such methods in the study of lexical and syntactic parasynonymy now represents an important line of research in Cognitive Linguistics. Amongst others, Fischer (2000), Schmid (2000), Gilquin (2003, 2006), Gries (2003, 2006), Grondelaers and Geeraerts (2003), Glynn (2004b, in press, forthcoming), Newman and Rice (2004, 2006), Heylen (2005), Dmitrieva (2005), Divjak (2006) Divjak and Gries (2006), Lemmens (2006), Wulff (2006, et al. 2007), Janda (2007), Frohning (2008), Gries and Divjak (this volume), Coleman (in press), and Janda and Solovyev (forthcoming), are representative of this movement. Both the syntagmatic and paradigmatic parameters are covered as well as conceptual-functional meaning and social-regional variation in usage.

However, extending this methodology to stand as a programmatic method in Cognitive Semantics faces three fundamental hurdles. These difficulties are analytical rather than theoretical in nature and their solutions lie in methodological development. The next section treats each problem in turn.

3. Frames, syntax, and the social dimension. A model for quantitative analysis

3.1 *Tertia Comparationis* and feature analysis

To date, no study has used such quantitative techniques to describe the polysemy of an abstract concept. This is due to the fact that a quantitative approach to polysemy necessitates a *tertium comparationis* (cf. Lewandowska-Tomaszczyk 1998, 1999; Glynn 2004a, 2006a). For abstract concepts, this is impossible since the designatum is a conceptual construct based in culture, possessing no *Lebenswelt* referent. Without an objective constant as a basis for semantic analysis, it is difficult to develop criteria that are sufficiently rigorous to permit the application of quantitative techniques. This problem of operationalising annotation for semantic features cannot be underestimated.

One solution to this problem lies in Frame Semantics (Fillmore 1985). Following Dirven et al. (1982), Rudzka-Ostyn (1989, 1995), and Fillmore and Atkins (2000), Glynn (2006b, 2006c) proposes a solution that goes some way to solving this problem. The reasoning is that one uses the semantic frame as the constant upon which feature studies may be based. This allows both the vagaries of situation context as well as the complexity of the cultural model and its encyclopaedic semantics to be handled in the analysis. What is more, the different arguments and their relations may be treated as semantic features in

the model. The drawback of this approach is twofold. The semantic frame must be posited *a priori* which leads to the same problems that have always plagued intensional definitions and, secondly, such an analytical model is obviously biased towards verbal forms and the concepts that are profiled by them.

Although there seems to be no way of resolving the first issue, this definitional procedure is “overt”. By positing a semantic frame and identifying the arguments and argument relations, the steps taken to define the concept are specified, thus verifiable and less susceptible to analytical flaws resulting from subjective bias. The second issue is more problematic. Although frame semantic structures are still valid for non-verbal concept profiling, if most of the frame arguments are maximally backgrounded and not linguistically expressed in the utterance, the model's ability to capture semantic structure is limited. Thus, for non-verbal profiles, further *ad hoc* parameters may need to be evoked in semantic analysis. This remains a weak point in the model.

3.2 Syntagmatic and paradigmatic dimensions of polysemy

Although both the syntagmatic and paradigmatic dimensions are covered in cognitive approaches to semasiological structure, the interaction between these structures is still not fully understood. The basic problem is that we have not established how schematic and/or morpho-syntactic semantics and less schematic lexical semantics interact. Within Cognitive Linguistics, one position is that syntactic semantics override or coerce lexical semantics (e.g. Talmy 2000). Another position is that there exists a complex interaction between all the various semantic structures in all degrees of schematicity (e.g. Langacker this volume). Following the research presented in Glynn (2002, 2004b, 2008), we make the assumption that syntactic variation affects a polysemy network, and that its effect cannot be satisfactorily predicted by positing meaning structure associated with grammatical forms and classes *a priori*. We must, therefore, account for this variable as an integral part of semantic description. The ramifications of this final point are important. It means that for a given lemma, or root lexeme, there will be semantic variation depending on its syntagmatic context. In other words, its collocation, grammatical class, and even tense or case will necessarily affect the meaning of the item.⁸

This may seem obvious, but to date, within Cognitive Linguistics, the role of this parameter in meaning description has not been considered. We adopt the solution presented in Glynn (in press): instead of treating the polysemy structure of a lexeme as it is expressed for a single part of speech, each lexeme is treated as a onomasiological field, or set of paronyms. This should allow the investigation to meet up with the current movement in Collostructional Analysis (Stefanowitsch and Gries 2003; Gries and Stefanowitsch 2004).

8. The term ‘lemma’ is used following the parlance of corpus linguistics, where it signifies the range of formal variants of the root lexeme.

3.3 Extra-linguistic variation and meaning as usage

Different people use different words in different situations in different ways. This, it would seem, is an undeniable fact of language. In a usage-based approach to language, we must necessarily account for the extra-linguistic factors that this simple statement entails. In other words, the usage and therefore meaning of a lexeme is different in different situations and this semantic variation, or polysemy, is our object of study. The methodology presented here is an attempt at developing a procedure that accounts for extra-linguistic factors, while it is in keeping with the theoretical tenets of Cognitive Linguistics. The basic assumption is that rather than identify conceptual structure, we identify the *various factors of usage that are a result of the conceptual structure that speakers associate with a given form*. The principal factors are dialectal (regional variation), sociolectal (social variation), and register-specific (medium variation).

It seems that the only way to describe the effect of such factors on usage is through corpus-driven quantitative research. The principle is simple: the co-occurrence of features, relative to given “factor variables”, represents structural tendencies in the use of a form. In other words, clusters of semantic features and/or the absence of features in given linguistic and extra-linguistic contexts, are indicators of the meaning of a word. Biber (1995), one of the most important figures in the application of this method, describes this succinctly. In the following quote, his “communicative functions” could be paraphrased as the conceptualisation associated with a form.

Factor interpretations depend on the assumption that linguistic co-occurrence patterns reflect underlying communicative functions. That is, particular sets of linguistic features co-occur frequently in texts because they serve related sets of communicative functions. (Biber 1995: 115)

In various forms, the work of Dirven et al. (1982), Geeraerts et al. (1994), Fischer (2000), and Schmid (2000) makes this assumption and it is accepted here *a priori*.

This assumption is most important and has strong implications for semantic research generally. Employing this method and accepting this assumption means that instead of positing senses and attempting to distinguish them, we simply identify patterns of usage. Degrees of distinctiveness between these patterns may be treated as a statistical question: relative to a given situation (referent, register, region etc.) what is the probability that a given pattern will be used? Approached in this manner, sense identification and distinction are merely summaries, albeit useful ones, of the multifactorial complexity of real language use.

4. A quantitative multifactorial case study of polysemy: *hassle*

In this section we follow a simple case study that shows how a Quantitative Multifactorial method reveals semantic structures that other methods cannot. We examine the semantic structure of the lexeme *hassle* in British and American English. The Fifth Edition of Shorter Oxford English Dictionary and Webster’s Third International Dictionary define *hassle* as:

verb trans. and intrans. Bother, pester, harass (a person); quarrel or wrangle *over* (something). *noun.* A problem, a difficulty; a quarrel, an argument; fuss, bother. (SOED)
n 1. heated argument : WRANGLE. B: a violent skirmish: FIGHT. 2: a protracted debate: CONTROVERSY. 3 a: a state of confusion or commotion: TURMOIL. B: a strenuous effort: STRUGGLE. *vi* ARGUE, FIGHT, DISPUTE. (Webster's)

Although a semantic analysis should not be compared with a dictionary entry, when people are asked to consult their intuition, these definitions seem adequate. There is no mention of variation between American and British, nor any mention of grammatical variation outside the two verbal forms and two nominal forms (the latter implied in the British definition) of *hassle*. Let us see what a coarse-grained quantitative usage-based and multifactorial investigation reveals.

4.1 The corpus and annotation

The mainstay of corpus-driven research focuses on syntactic structures and to these ends powerful parsing technology exists. However, for lexical semantics, tagged corpora are less essential. The most important features of a corpus for the study of content words are its representivity and sheer size. Content words repeat infrequently and lexical variation is typically sensitive to extra-linguistic factors. These two conditions mean that for a lexical semantic study to capture any degree of semantic subtlety of even the most common usages associated with a given lexeme, the corpus must be large and preferably representative of various types of language and register. It is for these reasons that commercially available corpora are less appropriate for lexical semantic investigations. One alternative is to use the internet as a source from which one may build a corpus. Although the internet suffers from thematic bias in its emphasis on “new technologies” as well as a bias of age and social class, the various media that it includes (such as Internet Relay Chat, Usenet, news press, blog-diaries, etc.), mean that a reasonable range of language types is represented.

The current study uses two corpora. The first was made using commercially available “web spiders” that allow one to download large quantities of internet files of a specific kind and from specified servers. The ability to select servers allows one to be reasonably certain about the origin of the text, which is important for concerns of dialect variation. The second corpus was developed by D. Speelman at the University of Leuven and is made up of data extracted from the LiveJournal on-line diary server. LiveJournal represents possibly the largest blog server currently in existence and it kindly allowed us to extract our text. One important feature of the LiveJournal database is that blog-writers must identify which secondary school they attended along with its address. This allows us to be almost entirely certain as to the dialectal origin of the text.

Despite a reasonable range of language types and topics of discourse, these corpora are not as representative as one would normally wish. In order to account for content bias, both theme (topic of discourse) and register (or language type) are systematically annotated. These two parameters as well as dialect, American English versus British English, make up our extra-linguistic factors. Since, within Cognitive Linguistics, we hold that language is a symbolic pairing of form and meaning, the examples are also

annotated for these two basic linguistic parameters. However, in order to operationalise the annotation and render the feature analysis as objective as possible, these parameters were further broken down into a range of variables. These variables can be summarised as four variable groups, two for the parameter of meaning and two for form. Firstly, for the parameter of form, the morpho-syntactic variables were coded separately from the argument structure. This means that the largely objective task of identifying part-of-speech, tense, and so forth can be kept separate from the more theoretically dependent criterion of argument structure.

For the semantic parameter of the feature analysis, the annotation is again divided. Firstly, the more objective of the two variable groups is that of argument types and their relations. Here information such as animacy versus inanimacy and abstractness versus concreteness, as well as (non) familiarity between actors or their power relations, and so forth, are annotated. This is largely objective, although for the adjectival and nominal profilings, this annotation becomes less insightful. However, for the instantiation of a semantic frame, we must note that although one may still annotate backgrounded participants by looking back in the text, this is not always practically possible. A distinction was maintained in the annotation between overt (i.e. profiled and linguistically expressed in the utterance) and covert (backgrounded and not expressed in the utterance) arguments. When there was doubt as to the nature of an argument or an argument relation, it was not annotated.

The second dimension, or variable group, of semantic annotation was the effect on the patient. It is similar to the stimulus feature in the FrameNet project. This feature is highly subjective and thus its results must be treated with caution. The variable includes twelve reasonably fine-grained distinctions that attempt to capture the “effect” upon the patient, such as a “request” being made of the patient or that the patient “feels imposed upon” or “interrupted”. In total, for both the formal and semantic features, 24 variables were annotated. Some of these, such as the morpho-syntactic tagging, were made up of more than 30 binomial values.

4.2 Techniques for Quantitative Analysis

Once the feature analysis is complete and all data are annotated, we need to search for correlations in the frequency of features, relative to the different variables. The field of statistics has an abundance of analytical techniques open for both exploratory investigation and hypothesis testing. The former is used to look for patterns in the data that may be informative; the second is used to determine if these patterns are significant or merely coincidental for a given dataset. This second step is essential since no dataset, no matter how large, can ever represent the reality of the population, in linguistic terms, the culmination of utterances that make up a language.

The results of the feature annotation take the form of cross-tabulations of frequencies, or contingency tables. This may be rephrased as: how often given features occur relative to the different variables specified. In order to examine these results statistically, we must choose from the wide range of exploratory techniques available for the study of categorical data. Our choice is determined by two factors. Firstly because manual semantic coding is labour intensive, our frequency results are relatively small. Obviously the larger the

dataset, the more reliable the results are, and therefore the more statistical techniques become reliable. Secondly, our data are categorical; they are made up of cross-tabulated frequencies of observed features. In other words, either feature x is present or it is not. However, many of the most powerful statistical techniques are designed for the treatment of continuous data and are not appropriate for our categorical results.

Categorical Principal Component Analysis is one of the possible techniques for the treatment of feature frequency in lexical analysis. However, this is a relatively new technique and has not yet been widely applied to this sort of data. Of the various cluster techniques, Model-Based Cluster Analysis may also be appropriate. This has the advantage of being suited to categorical data but has the disadvantage that the number of clusters must be specified before the analysis. Hierarchical Cluster Analysis is another option and is successfully used by Rice (1999), Divjak (2006, and Gries 2006), Gries (2006), Gries and Divjak (this volume), and Gries and Stefanowitsch (in press). Since we are exploring methodological techniques, it would be useful to examine a different method. The method employed here will be Correspondence Analysis. This technique is amongst the simplest to apply and is suited to categorical data.

All these techniques are designed for exploratory analysis. They are used to look for patterns in the data that *may* be representative of significant structure. However, they do not estimate the probability that a given correlation is statistically significant. Significance is the likelihood that the correlation, or relationship between sets of features and variables, is representative of the language as a whole and not just a coincidence in the dataset. In other words, we need to test to determine the probability that a given pattern observed in the data is a result of real factors and not merely coincidence. Obviously, the smaller the number of examples, the harder it is to be sure that the results are representative of the complexity of language reality.

There are many mathematical tests one may use to determine “statistical significance”. There exist also predictive techniques that examine many different variables simultaneously and even offer information as to the relative importance, or effect, of the different variables on the data. Logistic Regression Analysis and Log-Linear Analysis are probably the most appropriate for semantic research. However, for our current purposes, Correspondence Analysis, combined with certain significance tests, should suffice.

4.3 Analysis. The interplay of formal, semantic, and extralinguistic variables

Firstly, let us begin with the formal variation of the lemma. Although it is no secret that different parts of speech or certain collocations result in semantic variation for a single lemma, this parameter is rarely accounted for in polysemy study. Typically one form is chosen and considered in isolation. This, of course, runs contrary to the tenets of Cognitive Linguistics where the different forms associated with the lexical category are choices available for the different profilings of that category. It follows that we should attempt to account for this variation.

The corpus reveals a range of adjectival, nominal, gerundive, and verbal forms of the lemma *hassle*. The attributive adjectival forms are relatively infrequent. Example (1) is typical:

- (1) Specifically, she is interested in how hassling events influence attitudes toward re-frequenting a particular store. <www.uncw.edu/aa/2005-2006/csb.html>

The predicative adjectives are common. Although two possible argument structures are possible, a simple stative and a second that expresses the cause of the ‘hassle’ with an oblique, the former is extremely rare. The vast majority of cases overtly express the Cause, which is introduced by *with*, *by*, *over*, *at*, *because*, *for*, or *due to*.

- (2) I keep an eye on things when I can but I’m well hassled by numerous stuff at the moment. <[news:9bp2e4\\$rj\\$1@pegasus.csx.cam.ac.uk](mailto:news:9bp2e4rj1@pegasus.csx.cam.ac.uk)>

For practical reasons, we will not examine the adjectival forms in any depth.

The nominal variation is important. Other than the gerund, there are mass nouns and both singular and plural count nouns. Let us look at some frequencies of this form relative to dialect. Firstly, we may conflate the singular and plural forms of the count noun. Using the Binomial Exact Test and the Proportional Chi-Squared Test, it is very improbable that there exists a significant difference relative to each other or relative to the dialect variation. In our comparison, we may also include gerunds. However, the gerundive examples pose certain problems in the annotation of their frame structure. Following the FrameNet project, one may divide the gerund examples into “verbal” and “nominal” examples. Although this may at first seem unnecessary, the examples clearly separate into instances where the gerund is part of an event structure and where it is part of a nominal profiling. Nevertheless, relative to dialect, both the nominal and verbal gerunds behave in the same manner, both being highly associated with American English.

In Table 1, we see that although there are relatively similar frequencies of nominals across the two dialects (179 out of 344 and 198 out of 347 occurrences in the respective dialects), their plexity is far from uniform. The Proportional Test is used to demonstrate that it is highly probable that these differences are significant. The *p*-values are listed in the table. Any figure less than 0.05 should be read as significant. Obviously, the closer to zero, the “more reliable” that degree of significance becomes. So, for example in Table 1, a *p*-value of 2.2e-16 (or 0.000000000000000022) is extremely significant. The difference in the frequency of the gerund is also significant, but less so. The Proportional Test uses the Chi-square algorithm and so becomes unreliable with figures under 10. However, the same test applied to the relative British and American frequencies of the verbal gerundive (UK: 4/161, US: 22/143) gives us a *p*-value of < 0.001. Finally, if the nominal and verbal gerundives are combined to give a purely formal category of “gerund” (UK: 6/344, US: 34/347), the test still gives a *p*-value of < 0.001. Therefore, we can be sure that the gerund is highly associated with American English and highly dissociated with British. The different

Table 1. Nominal frequency relative to dialect

Nominal variation	UK	US	Total	Prop. Test
Mass	144	70	214	2.2e-16
Count	33	116	149	1.30e-12
Nom. gerund	2	12	14	0.02199
Total	179	198	377	

grammatical profiling of mass noun, count noun, and gerund obviously represent different meanings of the lexeme. This is an example of the effect of an extralinguistic factor upon semantic structure. Such extralinguistic concerns cannot be sidelined as question of social variation. They make up part of our understanding of how language is used, and therefore part of the encyclopaedic semantics associated with a given form. Let us turn now to the interplay of formal and semantic factors.

In an effort to capture semantic structure without direct annotation of semantic features that tend to be subjectively determined, we can draw on the model of Frame Semantics and annotate Actor types and Relations. In order to see the difference in usage between the nominal and the verbal examples, we may examine the differences in which Actor types are associated with which grammatical constructions. Several of the most common Actor types include a known or “specified” human (Hum_spec), an unknown or non-specified human (Hum_NtSpec), abstract events (Ab_Evnt), concrete events (CcrT_Evnt), abstract states-of-affairs (Ab_SoA), and concrete things (Thing). Using Correspondence Analysis, we may investigate the relationship between these different Actor types and the constructions. Four constructions are considered here: simple transitive verbs (Trans), transitive verbs with an oblique argument (Trans-obl), resultatives (Result), and nominals (Nominal). Correspondence Analysis uses a reasonably simple statistical technique to examine relative degrees of association. The resulting plots are should be interpreted visually, correlations being depicted by relative proximity. The numbers indicated on the axes are there to help determine this relative proximity.

The first obvious grouping (i) is the association between the nominals and inanimate actors. Here, Cause-Actors that are abstract and concrete events, things, and states-of-affairs are clustered with the nominal profiling. This is contrasted with group (ii) that brings together the three verbal constructions and the animate Cause-Actors types. Although an institution is technically not animate, it is clearly a borderline case. In the data, institutions

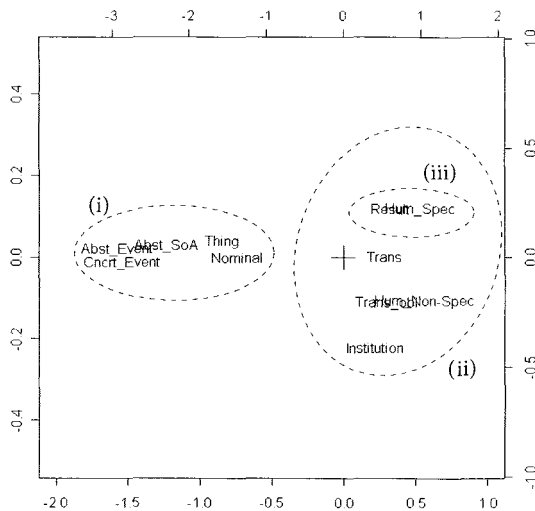


Figure 1. Correlation between Parts of Speech and Actor types

were often whole-for-part metonyms where a bank, university, or some other institution was used to stand for the individuals ‘hassling’ the Patient.

- (3) Debtors have the right to ask collection agencies or any source hassling them for debt collection to stop. <www.public.asu.edu/~hkartadi/laws_in_credit_repair>

More specifically, for the third cluster, the resultative constructions are highly associated with “specified human” Actors. Indeed, they overlap to the extent that the plot is difficult to read. The simple Transitive argument structure is neutral in regard to the specificity of the human feature where the Transitive-Oblique construction is associated with non-specified humans and also institutions. Notice also that known specified human through to institution is depicted as a cline from top to bottom of the group. This suggests that animacy is indeed related to the different Argument structures and grammatical constructions, the resultatives being highly associated with animate Cause-Actors but that this animacy becomes less important for the Cause-Actors of simple Transitive constructions and then less again for the Transitive-Oblique constructions. Let us consider some examples:

- (4) a. since I still have to go through all the hassle of US Immigration
<forum.flightmapping.com/forum_posts.asp?TID=591andget=last>
b. ... my physics exam tomorrow afternoon ... I have been switching rapidly between thinking that it’s going to be a piece of piss and thinking I should just top myself now and save the hassle.
- (5) a. Well, I had been hassling Argo for a while to implement my ‘wish list’ into ...
www.heyrick.co.uk/voyager/newsagent/intro.html
b. You were the one hassling me for an answer.
<www.thefridayproject.co.uk/talk/archive/index.php/t-710-p-2.htm>
- (6) a. Some smokers also have a dream that someday the non-smoking world will quit hassling them about their smoking.
<www.nap.edu/books/0309064090/html/169.html>
b. It had me chortling for minutes. If the money’s in his wife’s account, why aren’t they hassling her? <www.thefridayproject.co.uk/talk/archive/index.php/t-2162.html>

Example (4) is typical of the examples captured by the plot in group (i). The choice of the nominal profiling for such events is quite “logical” and we can safely say that non-animate Cause types are suited to the “meaning” expressed by the nominal profiling. Similarly, example (5) represents what seems to be intuitively clear. It is intuitively sound that known human Causes should be common in the resultative examples. What is perhaps less obvious is the animacy cline from familiarity through unfamiliarity to institutional Cause Actors correlating with resultative, simple transitive, and transitive-oblique constructions. This is visible if we compare examples (5) with those in (6).

However, with small frequencies, Correspondence Analysis becomes sensitive to distortions and is less reliable. Although the fact that Transitive-Oblique constructions are highly associated with unfamiliar-human Cause-Actors is intuitively reasonable, it needs further corroboration. It is crucial to remember that this technique is merely an exploratory technique restricted to positing *possible* linguistic structure. Example (6) is offered to show the type of examples that the rather heterogeneous group (iii) represents.

The Intransitive construction was not added to the above Correspondence Analysis because it had low frequencies for all but one of the Actor types. If we look at a couple of the Intransitive examples, we will see why this is a distinct usage.

- (7) a. Officer McCoy, me and him was hassling and my gun went off...
 <www.privacy-council.org.uk/files/other/forrester%20Bowe-rtf.rtf>
 b. It made all the surfers really spread out and we weren't all hassling on the one peak.
 <oneillcwc.asglive.com/daysix>

As we see in the example (7), this usage is semantically distinct from examples (4) to (6). The Intransitive construction is relatively infrequent in the corpus and is semantically marked (despite its prominence in the dictionary entries cited above). This construction draws our attention to an inherent weakness in the quantitative method: infrequent occurrences cannot easily be taken into account. This is a reminder that a corpus-driven quantitative investigation should work in tandem with other methods. The markedness of the Intransitive brings us to the importance of the different constructions associated with the verbal form of the lexeme.

Six basic syntactic forms are revealed. These constructions fall out from the annotation of the different arguments as various semantic roles. In order to capture the different argument structures, the semantic roles of Actor, Cause, Patient, Instrument, and Goal were employed. These were assigned to Subject, Object, and Oblique for all verbal examples. Let us examine the constructions that result from this annotation.

- A. Cause-Oblique Transitive
 Sub. Pred. Obj. Obl.
 She hassled me because of my spots.
 Act. Pred. Pat. Cause.
 ex.: We should all quit hassling the nice oil companies about profits.
 <www.alternate-heaven.com>
- B. Transitive
 Sub. Pred. Obj.
 He hassled me.
 Act./Cause. Pred. Pat.
 ex.: In fact, she hassled him so thoroughly on the street that she made him...
 <www.izzlepaff.com/blog/archives/2004/03/>
- C. Resultative
 Sub. Pred. Obj. Obl.
 She hassled me to eat.
 Act./Ag. Pred. Pat. Goal.
 ex.: If you are caught without the necessary papers/stamp, they hassle you into thinking they will detain... <www.brama.com/travel/messages/4994.html>
- D. Subject-Patient Transitive
 Sub. Pred. Oblique
 He hassled over it.
 Act./Pat. Pred. Cause.
 ex.: Everyone these days is hassling over their weight when they should be hassling about their families. <www.faqfarm.com/Q/How_long_does_it_take_a_12-year-old_to_lose_weight>

- E. Instrumental-Oblique Transitive
 Sub. Pred. Obj. *with*-Oblique
 He can hassle you with spies.
 Act. Pred. Pat. Ag.
 ex.: Stop hassling me main man with constant calls! <www.o2.co.uk/services/messaging/voicemail901/celebrityvoicemail>
- F. Intransitive
 Sub. Pred.
 She hassles
 Act. Pred.
 ex.: see example (6).

There is, of course, considerably more formal variation than this, especially in the various types of oblique. Differences between *hassle into* something and *hassle to do* something or *hassle over* and *hassle with* represent more than synonymous formal variation. Such variation has clear semantic characteristics. For example, the Cause-Oblique Transitive construction (A) combines with a very wide range of oblique forms, coded with prepositions such as *for*, *because*, *over*, and *on*. It is most likely that such syntactic variation contributes to the polysemy of the lexeme. Here we must, however, restrict ourselves to a coarse-grained investigation.

Firstly, let us see if these forms occur equally across the dialects considered. Again we can use the proportional test to determine whether the differences between the frequencies of each construction are significant.

The results of the proportional test reveal a significant difference between the dialects in the use of the Patient-Oblique Transitive and Resultative Constructions. The latter is especially associated with British English and not American. The other clear difference between the dialects is the use of the Patient-Oblique Construction, which is rare in British but relatively common in American.

Since we have seen that the two dialects are markedly different in their constructional variation for the lexeme, we should investigate how those constructions are associated with different senses for each dialect separately. Once again we can turn to the exploratory technique of Correspondence Analysis for a visualisation of the correlations. If we combine the frequency tables of the constructional variation with one of the semantic variables, we may find correlations between these two variables. The variable of stimulus, or the kind of effect the Cause of the hassle-event has upon the Patient, must be annotated using subjective

Table 2. Dialect variation of construction

Construction	UK	US	Prop. Test
A Cause-Oblique Trans.	10	21	0.02463
B Transitive	110	81	–
C Resultative	34	12	0.003387
D Patient-Oblique Trans.	2	22	1.36e-02
E Instr.-Oblique Trans.	1	4	0.2996
F Intransitive	4	3	–

judgement. However, if results gleaned from its annotation are statistically significant and intuitively plausible, then we may tentatively employ this variable in our analysis.

The plot below represents the results of a Correspondence Analysis of six constructions against six different stimuli. The stimuli considered here are “interrupt-disturb” (interrupt), “repetition-boredom” (rep), “energy-time” (energy), “mental-stress” (mental), “imposition-intrusion” (impose), and “request-solicit” (request). The Intransitive Construction is not included since it was not annotated for stimulus.

Immediately, three correlations are apparent. Firstly, the proximity of the “request” stimulus to the Resultative Construction (C), indicated by (i) on the plot, is a result of what is an intuitively reasonable association. A large percentage of the resultative occurrences describe situations where something is being asked of the patient. Examples (8a)–(8c) are typical of the Resultative Construction–“request” feature correlation.

- (8) a. Be prepared it is pricey and if you’ve booked a table you’ll be given a maximum of 2 hours before the staff start hassling you to have desserts and/or last...
<www.viewlondon.co.uk/info_Pubbar_6564.html>
- b. You were the one hassling me for an answer, not the other way round
<www.thefridayproject.co.uk/talk/archive/>
- c. Chad Holleman was hassling Preston for the phone number of some girl he ...
<www.cheddarheads.co.uk/nfle/nfle01/week07.htm>

The relative frequency and importance of this association is clear. Consider below the frequency of occurrences of this construction combined with the “request” feature relative to other semantic features and other constructions. Although by no means a unique correlation, its association is relatively high. Below, we see how its correlation compares with other construction–stimulus feature cross-tabulations.




Figure 2. Construction–stimulus correlation in British English

Table 3. Request–resultative correlation in British English

Stimulus feature	Resultative Cx (C) (34 total)	Transitive Cx (B) (110 total)	Cause-Obliq. Cx (A) (10 total)
Request	21	25	1
Impose	18	86	8

The association between the “request” feature and the Resultative Construction relative to the other stimulus features and constructions should be evident.

Secondly, the association depicted in Figure 2 between the stimulus feature of “imposition” with the Transitive construction (B) and the Cause-Oblique Transitive construction (A) is no surprise. These two constructions are the most commonly occurring just as this stimulus feature, “impose”, is the most common. Moreover, this is true for both dialects. It seems plausible that the two basic transitive constructions should be associated with instances of people imposing themselves upon other people. An intuitive reading of the dataset would surely lead a non-quantitative linguistic analysis to propose this sense as a basic sense of verbal *hassle*. Examples (9a)–(9c) are typical.

- (9) a ... aggressive beggars who profit by hassling members of the public.
<archive.thisisoxfordshire.co.uk/2001/8/15>
b. I hate those charities who employ people to hassle you on the streets.
<www.tiscali.co.uk/forums/showthread.php?t=108045>
c. However, get there early and there'll be no door staff and no one will hassle you.
<www.edinburghmetalscene.co.uk/archive/index.php/t-12359.html>

Correspondence Analysis visualises the intuitively sound generality of this sense and identifies its correlation in a quantitative and automated fashion.

The third correlation (iii) is between the stimulus feature of “energy” and the Patient-Oblique Construction (D: *hassle over Pat.*). Due to the low frequency of this construction in British, we must be cautious in reading the plot at this point. However, again this correlation is intuitively valid.

- (10) a. Instead of hassling with multiple drives attached to different computers, you can back it all up to one central location – automatically.
<www.pixmania.co.uk/uk/uk/183689/art/maxtor/onetouch-iii-shared-stora>
b. No more hassling with your laptops touchpad or pointing stick.
<computing.kelkoo.co.uk/>

These examples are typical of those in question and they support the hypothesis that the meaning of the construction fits with the semantics of exhorting energy over something.

Let us turn to the American case. The raw results differ considerably. Let us begin by introducing another three stimulus features that may be relevant. In American English there is a reasonable number of examples where the act of judging the patient negatively is the cause of the state of hassle. These examples were coded as “condemnation”. The two other relevant stimulus features may be referred to as “repetition”, where the patient is hassled due to some repetitive event and “interruption”, where it is an interruption event that has caused the patient hassle.

Table 4. Dialect variation for stimulus feature

Stimulus	British	American	Prop. Test
request	47	23	0.01008
imposition	114	87	0.087
repetition	10	8	-
interruption	18	9	0.1961
energy	6	28	0.000679
condemnation	2	23	7.05e-03
thought	23	30	-

Notice the significant differences between the two dialects on the frequency of the stimulus features of “energy”, “condemnation”, and “request”. “Energy” and “condemnation” are significantly associated with American and not British, contrary to “imposition”, “request”, and “interruption”. Let us plot just the American results here against the constructional variation, once again using Correspondence Analysis.

Again, we see the grouping (i) that is a result of the association between the feature “energy” and Patient-Oblique construction (D). This further verifies the association witnessed for the British data, especially since the data are less sparse for this construction in the American dataset. Similarly, although less common in American, the semantic feature “request” is again highly associated with the resultative construction (ii). Lastly, and also similar to the British data, we have the rough grouping of the more common semantic features and more “basic” grammatical constructions (iii). That is to say, relative to the feature “energy” and the Patient-Oblique construction, as well certain other outliers, these more general semantic features and constructions are clustered. The principal difference between the British and American results here is the outliers. The Intransitive Construction (F) and perhaps also Instrumental-Oblique Transitive construction (E) seem to lack




Figure 3. Construction-stimulus correlation in American English

strong associations with any of these semantic features. Likewise, the semantic feature of “thought” is not associated with anything, though it is clearly dissociated from “request” and “condemnation”. This differs from British, where “repetition” was the outlier. Perhaps more data is needed to capture these relations or perhaps they are not particularly associated with any grammatical constructions or any other stimulus features.

Finally, not only does the similarity between the plots in Figure 2 and Figure 3 tell us that the correlations between this semantic variable and the grammatical constructions are relatively similar across the two dialects, having two different datasets serves as a test, adding weight to the hypothesis that these correlations are valid. The next logical step at this point would be to use the dialect as dependent variable and submit the data to a Logistic Regression Analysis. This technique is a confirmatory technique that would allow us to verify this hypothesis. We will not, however, move to confirmatory techniques in this study.

Let us examine another semantic feature. Each example was coded as to whether there was an element of humour involved. Again this feature is subjective in nature and so only very significant differences should be considered reliable. The results here show again a significant difference between the dialects as well as a strong association with one of the constructions.

To begin with, the use of humour and this lexeme is essentially a British characteristic. Out of the 161 British verbal examples, at least 72 were humorous against 17 out of 141 American examples. The proportional test gives a p -value < 0.001 , which should be significant enough to make up for almost any degree of error in the subjective nature of this annotation used. We will, therefore, focus exclusively on the British data for the description of this feature.

We see here that although humour has a high association with both the simple Transitive and Resultative constructions, its association with the Resultative is very significant relative to the overall number of Resultative examples. In other words, although there are more examples of humorous + Transitive co-occurrences, nearly all the Resultatives were humorous.

It must also be remembered that this construction is highly associated with the British dialect as well as the semantic feature of “request”. This was visible in Table 4, above. This shows how the “request” feature is associated with British relative to American, but also that “condemnation” is highly associated with American and not British. Now consider the relations between “humour” and the stimulus features for British.

Table 5. Construction–humour correlation in British English

Construction	+Humour	-Humour	Prop. Test
A Cause-Oblique Trans.	2	8	–
B Transitive	41	69	0.0534
C Resultative	26	8	2.01e-02
D Patient-Oblique Trans.	0	2	–
E Instrumt.-Oblique Trans.	0	1	–
F Intransitive	0	4	–
Total	69	92	

Table 6. Stimulus–humour correlation in British English


Stimulus feature	+Humour	-Humour
request	31	16
imposition	44	70
repetition	5	5
interruption	11	7
energy	1	4

Here we see a strikingly similar correlation to the correlation between the “request” feature and the Resultative construction relative to the “imposition” and the more semantically schematic Transitive construction presented in Table 5. Although there are more humorous examples that possess the stimulus feature of “imposition”, there are almost two times more humorous than non-humorous examples that possess the “request” feature.

This clustering of correlations is beginning to give the picture of semantic structure. The “request” stimulus is associated with the “resultative” construction and both with the humour feature, all three of which are typical of British English and not American. Many would call this a meaning of *hassle*.

We can subject these data to another Correspondence Analysis in order to visualise this correlation. Since this correlation is between three variables, we employ Multiple Correspondence Analysis which follows quite a different procedure but one that is based on similar mathematical principles and whose plots should be interpreted in a similar manner. Figure 4 plots the combinations of the frequencies of constructions and stimulus features against the occurrence and non-occurrence of the “humour” feature.

Firstly, the plot reveals the correlation of the Resultative construction (AS.C), the stimulus feature of “request” (stim.req), and the humour feature. However, the “humour” feature is, as we know from Table 5, also correlated with the simple Transitive construc-

**Figure 4.** Construction, stimulus, and humour correlation

tion (AS.B) and the “imposition” feature, cf. Table 3. We see this in the proximity of the two groups and the fact that the “humour” feature (hmr.Hmr) is also plotted close to the “imposition” feature and the Transitive. It is for this reason that the groups seem to overlap. The second grouping, of the “imposition” feature (stim.impos), the “-humour” feature (hmr.NtHmr), and the simple Transitive Construction, depicts the kind of usages we saw in example (9), arguably one of the most basic usages of the lexeme.

Importantly, one should note the distance between the features “repetition” (stim.rep) and “interruption” (stim.interrupt) and the constructions Oblique-Cause (AS.A) and Resultative (AS.C). Although we cannot say these stimulus features are highly associated with the simple Transitive construction (AS.B), relative to the other two constructions, there seems to be some association. This is one of the difficulties of reading plots produced in Correspondence Analysis. We cannot say, for instance, that “imposition” is more closely associated with the Oblique-Cause construction than “repetition” or “interruption”, because these latter two are also plotted in association with the Resultative and Transitive constructions. When we look at the frequencies for these correlations this becomes clear. In fact, these two stimulus features had zero co-occurrence with both the Transitive and Resultative constructions, yet 11 (repetition) and 17 (interruption) occurrences with the Oblique-Cause construction. Although small numbers, relative to the 161 British verbal occurrences, this is not negligible.

5. Summary

Although nominals are equally distributed across the two dialects, closer inspection reveals variation. In American, these grammatical classes are generally profiled by gerunds and count-nouns; where in British one finds almost exclusively mass-noun profiling. The semantic variation this almost surely entails was not investigated. However, semantic variation was investigated between the nominal and verbal profilings. A Correspondence Analysis revealed the different constructions associated with the lemma and correlated them with Agent types. This points to what seems to be a clear tendency: the nominal forms are highly associated with Agents that are either events or inanimate things. These Agent types constitute a cline from events through inanimate things to familiar-humans. Verbal forms are at the other end of this continuum and are highly associated with familiar-human Agents. We can describe this difference in usage as polysemic structure. The same analysis discerned another pattern. At the animate end of the continuum, a sub-cline appears, tending from familiar-human through unfamiliar-human to institution. Along this cline of Agent types, the Resultative construction tends to be associated with familiar-humans in contrast to the Transitive and Transitive-Oblique constructions which tend towards “less” animate Agents.

Focusing on verbials, the annotation revealed that six grammatical constructions are associated with the verbal forms of the lemma. Instead of performing a Collostructional Analysis, we continued to investigate the different semantic features associated with each lexeme-construction pairing relative to dialect. Firstly, some construction-pairs seem highly associated with one of the two dialects. For example, the Resultative is essentially

a British usage and the Passive-Oblique Transitive is almost exclusively American. Due to this formal variation between the dialects, the semantic structure was examined for each language variety independently. One of the constructions, the Intransitive construction, was semantically distinct, but its frequency too low to be included a quantitative study.

The semantic analysis focused on two semantic variables. The first of these, the 'stimulus' feature, correlates with different constructions in significantly different ways. However, despite the 'stimulus' feature variation between constructions and the variation in construction frequency between the dialects, the stimulus-construction associations behave in a similar manner across the dialects. In both cases, the Resultative construction is highly associated with the stimulus of "request" and, importantly, highly disassociated with the other stimulus features and constructions. This adds weight to the first Correspondence Analysis, which suggests that the Resultative is distinct in its association with familiar-human Agents. In addition, a second distinct association was revealed. The Subject-Patient Transitive construction is highly and distinctly associated with the stimulus feature of "energy". Again this is true of both dialects. Regrettably, this could not be treated in depth. A final correlation brought out by these two Correspondence Analyses is that for both dialects there is a clear association between Cause-Oblique Transitive and the stimulus of "imposition", where the patient feels put upon by someone or something. Unfortunately, this association also had to be left aside for future investigation.

Finally, we focused on the typically British usage that combines the Resultative construction, the stimulus of "request", and the familiar-human Agent. To this, we added the second semantic variable, "humour". Here again, we see a highly significant association with the British usage of the lexeme. Moreover, there is strong correlation between the "humour" feature and the "request" feature, as well as between "humour" and the Resultative construction. The grouping of these features, all of which are distinct relative to other features and forms and specific to British English also coincides with the association of familiar-human Agents that was revealed in the first Correspondence Analysis. Together such a cluster of forms and usages could be reasonably argued to represent a "meaning" of *hassle*.

This possibility was considered in a Multiple Correspondence Analysis. Its results corroborate what is seen in the individual cross-tabulations. However, the analysis suggests that although humour is highly associated with the Resultative-"request" correlation, it is also somewhat characteristic, at least in British, of another more general meaning. In that, the most common construction, the simple Transitive construction, correlates with the most common stimulus feature, "imposition" and this form-meaning pair lies in between the "humour" and "non-humour" features in the results of the Multiple Correspondence Analysis. This finding is intuitively reasonable since the use of humour is generally characteristic of the British dataset.

This raises a final important point. Although these correlations of semantic features such as humour and linguistic forms such as the Resultative can be argued to represent part of the polysemic (semasiological) structure of *hassle*, we must be wary of thematic bias. In other words, these characteristics might be features of the corpus rather than the lexeme in question. Since the kind of language found on the Internet tends to belong to younger speakers and is biased for their topics of discourse, it is possible that this misrep-

resents the importance of this usage in British English. However, the two datasets, British and American, are of the same text type, so we can confidently say that for the quasi-spoken language of the blog-diaries, the differences between the two dialects, relative to this language type, are valid generalisations. Nevertheless, it seems reasonable to assume these findings do represent the language as a whole, but cohort studies with different language types need to be undertaken to verify this.

Last but not least, Correspondence Analysis is only an exploratory method and we are working with relatively small frequencies. The next step is to obtain larger frequencies and attempt to validate the observed patterns using statistical confirmatory techniques such as Log-Linear Analysis and Logistic Regression Analysis. The point of the study was to test the methodology and show that it may reveal semantic structure not detected through intuition-based analyses. In this, the study has proved successful and the integration of results from this kind of investigation with results obtained through a Principled Polysemy analysis should be a straightforward endeavour. Future work needs to operationalise the integration of the Quantitative Multifactorial method and the Principled Polysemy Model for the study of sense variation.

Cognitive Linguistics is a usage-based theory of language and one that assumes language is driven by our encyclopaedic knowledge of the world. In light of this, the kind of usage patterns that Quantitative Multifactorial methods identify offer important clues to the conceptual structures associated with linguistic forms. Although, presenting the results in terms that are typical of the cognitive research community still needs development, mapping the usage, and therefore meaning, of lexemes and constructions is precisely in keeping with the lexical semantic tradition developed by Lakoff (1987). The principal difference is that such quantitative results offer relative tendencies rather than 'different meanings'. This, however, seeing the complex and varied nature of language, is arguably a more cognitively realistic approach to the description of the conceptual structure.

References

- Atkins, B. 1994. Analyzing the verbs of seeing: A frame semantics approach to corpus lexicography. *Proceedings of the Twentieth Annual Meeting of the Berkeley Linguistics Society*: 42–56.
- Biber, D. 1995. *Dimensions of Register Variation*. Cambridge: Cambridge University Press.
- Bellavia, E. 1996. The German *über*. In *The Construal of Space in Language and Thought*, M. Pütz & R. Dirven (eds), 73–107. Berlin: Mouton de Gruyter.
- Boers, F. 1996. *Spatial Prepositions and Metaphor: A Cognitive Semantic Journey Along the Up-Down and Front-Back Dimensions*. Tübingen: Narr.
- Brugman, C. 1983. *The Story of Over: Polysemy, Semantics, and the Structure of the Lexicon*. Trier: LAUT.
- Colleman, T. In Press. Beyond the dative alternation: the semantics of the Dutch *aan*-dative. In *Usage-Based Cognitive Semantics. Corpus-Driven Approaches to Conceptual Structure*, D. Glynn & K. Fischer (eds). Berlin: Mouton de Gruyter.
- Croft, W. 1998. Linguistic evidence and mental representations. *Cognitive Linguistics* 9: 151–173.
- Croft, W. This Volume. Toward a social Cognitive Linguistics. In *New Directions in Cognitive Linguistics*, V. Evans & S. Pourcel (eds). Amsterdam: John Benjamins.

- Cruse, A. 1995. Polysemy and related phenomena from a cognitive linguistic viewpoint. In *Computational Lexical Semantics*, P. Saint-Dizier & E. Viegas (eds), 33–49. Cambridge: Cambridge University Press.
- Cuyckens, H. 1993. The Dutch spatial preposition “in”: A cognitive-semantic analysis. In *The Semantics of Prepositions. From Mental Processing to Natural Language Processing*, C. Zelinsky-Wibbelt (ed.), 27–72. Berlin: Mouton de Gruyter.
- Cuyckens, H. 1994. Family resemblance in the Dutch spatial preposition *op*. In *Kognitive Semantik. Ergebnisse, probleme, Perspektiven*, M. Schwartz (ed.), 179–196. Tübingen: Narr.
- Cuyckens, H. 1995. Family resemblance in the Dutch spatial prepositions *Door* and *Langs*. *Cognitive Linguistics* 6: 183–207.
- Cuyckens, H., Sandra, D. & Rice, S. 1997. Towards an empirical lexical semantics. In *Human Contact through Language and Linguistics*, B. Smieja & M. Tasch (eds), 35–54. Frankfurt: Lang.
- Deane, P. 1988. Polysemy and cognition. *Lingua* 75: 325–361.
- Deane, P. 2006. Multimodal spatial representation. On the semantic unity of *over*. In *From Perception to Meaning. Image Schemas in Cognitive Linguistics*, B. Hampe (ed.), 235–284. Berlin: Mouton de Gruyter.
- Dirven, R. 1981. Metaphors of spatial relations. In *Anglistentag Trier 1981*, J. Hasler (ed.), 63–91. Frankfurt: Lang.
- Dirven, R., Goosens, L., Putseys, Y. & Vorlat, E. 1982. *The Scene of Linguistic Action and its Perspectivization by SPEAK, TALK, SAY, and TELL*. Amsterdam: John Benjamins.
- Divjak, D. 2006. Ways of intending: Delineating and structuring near-synonyms. In *Corpora in Cognitive Linguistics. Corpus-Based Approaches to Syntax and Lexis*, St. Th. Gries & A. Stefanowitsch (eds), 19–56. Berlin: Mouton de Gruyter.
- Divjak, D. & Gries, St. Th. 2006. Ways of trying in Russian: Clustering behavioral profiles. *Corpus Linguistics and Linguistic Theory* 2: 23–60.
- Dmitrieva, E. 2005. Об экспансии в русском языке констукций с предлогом *no*. *Scando-Slavic* 51: 67–83.
- Evans, V. 2004. *The Structure of Time: Language, Meaning and Temporal Cognition*. Amsterdam: John Benjamins.
- Evans, V. 2005. The meaning of *time*: Polysemy, the lexicon and conceptual structure. *Journal of Linguistics* 41: 33–75.
- Evans, V. & Tyler, A. 2004a. Rethinking English ‘prepositions of movement’: The case of *to* and *through*. In *Adpositions of Movement*, H. Cuyckens, W. De Mulder & T. Mortelmans (eds), 247–270. Amsterdam: John Benjamins.
- Evans, V. & Tyler, A. 2004b. Spatial experience, lexical structure, and motivation: The case of *in*. In *Linguistics Studies in Motivation*, G. Radden & K.-U. Panther (eds), 157–192. Berlin: Mouton de Gruyter.
- Fillmore, C. 1985. Frames and the semantics of understanding. *Quaderni di Semantica* 6: 222–254.
- Fillmore, C. & Atkins, B. 2000. Describing polysemy: The case of ‘crawl!’. In *Polysemy: Theoretical and Computational Approaches*, Y. Ravin & C. Leacock (eds), 91–110. Oxford: Oxford University Press.
- Fischer, K. 2000. *From Cognitive Semantics to Lexical Pragmatics. The Functional Polysemy of Discourse Particles*. Berlin: Mouton de Gruyter.
- Frohning, D. 2008. From grammatical variation to cognition: Insights from a case study on causal markers in German. In *From Gram to Mind: Grammar as Cognition*, J. R. Lapaire, G. Desagulier, & J. B. Guignard (eds), 167–188. Bordeaux: Presses universitaires de Bordeaux.
- Fuchs, C. 1987. Ambigüité, vague, polysémie, et continu. *Quaderni di Semantica* 8: 299–305.
- Fuchs, C. 1991. Polysémie, interprétation et typicalité. In *Sémantique et Cognition: Catégories, Prototypes, Typicalité*, D. Dubois (ed.), 161–170. Paris: CNRS.
- Geeraerts, D. 1989. Prospects and problems of prototype theory. In *Prospects and Problems of Prototype Theory*, D. Geeraerts (ed.), 587–612. [Linguistics 27]. Berlin: Mouton de Gruyter.
- Geeraerts, D. 1992. The semantic structure of Dutch *over*. *Leuvense Bijdragen* 81: 205–230.

- Geeraerts, D. 1993. Vagueness's puzzles, polysemy's vagaries. *Cognitive Linguistics* 4: 223–272.
- Geeraerts, D. 1997. *Diachronic Prototype Semantics. A Contribution to Historical Lexicology*. Oxford: Oxford University Press.
- Geeraerts, D. 1999. Beer and semantics. In *Issues in Cognitive Linguistics*, L. De Stadler & C. Eylich (eds), 35–55. Berlin: Mouton de Gruyter.
- Geeraerts, D. 2000. Saliency phenomena in the lexicon: A typology. In *Meaning and Cognition*, L. Albertazzi (ed.), 79–101. Amsterdam: John Benjamins.
- Geeraerts, D. 2005. Lectal data and empirical variation in Cognitive Linguistics. In *Cognitive Linguistics. Internal Dynamics and Interdisciplinary Interactions*, F. Mendoza Ibáñez & M. Peña Cervel (eds), 163–189. Berlin: Mouton de Gruyter.
- Geeraerts, D. 2006. Methodology in Cognitive Linguistics. In *Cognitive Linguistics. Current Applications and Future Perspectives*, G. Kristiansen, M. Achard, R. Dirven & F. J. Ruiz de Mendoza Ibáñez (eds), 3–31. Berlin: Mouton de Gruyter.
- Geeraerts, D., Grondelaers, S. & Bakema, P. 1994. *The Structure of Lexical Variation. Meaning, Naming, and Context*. Berlin: Mouton de Gruyter.
- Gilquin, G. 2003. Causative *get* and *have*. So close, So different. *Journal of English Linguistics* 31: 125–148.
- Gilquin, G. 2006. The place of prototypicality in corpus linguistics: Causation in the hot seat. In *Corpora in Cognitive Linguistics. Corpus-Based Approaches to Syntax and Lexis*, St. Th. Gries & A. Stefanowitsch (eds), 159–192. Berlin: Mouton de Gruyter.
- Glynn, D. 2002. LOVE and ANGER. The grammatical structure of conceptual metaphors. *Cognitive Approaches to Metaphor*. Special edition of *Style* 36: 541–559.
- Glynn, D. 2003. Polysemy and over-polysemy. A three tiered approach to sense variation. Paper presented at the 8th International Cognitive Linguistics Conference. Logroño, Spain.
- Glynn, D. 2004a. Cognitive constructs. Perceptual processing and conceptual categories between mind, language, and culture. In *Language, Culture, and Cognition*, A. Soares da Silva, A. Torres & M. Gonçalves (eds), 247–268. Coimbra: Almedina.
- Glynn, D. 2004b. Constructions at the crossroads. The place of construction grammar between field and frame. *Annual Review of Cognitive Linguistics* 2: 197–233.
- Glynn, D. 2006a. Iconicity and the grammar-lexis interface. In *Insistent Images. Iconicity in Language and Literature*, E. Tabakowska & C. Ljungberg (eds), 267–286. Amsterdam: John Benjamins.
- Glynn, D. 2006b. Concept delimitation and pragmatic implicature. *Perspectives on Metonymy*, K. Kosecki (ed.), 157–174. Frankfurt: Lang.
- Glynn, D. 2006c. Conceptual metonymy – A study in cognitive models, reference-points, and domain boundaries. *Poznań Studies in Contemporary Linguistics* 42: 85–102.
- Glynn, D. 2008. Arbitrary structure, cognitive grammar, and the *partes orationis*. A study in Polish paradigms. In *Naturalness and Iconicity in Linguistics*, K. Willems & L. De Cuypere (eds), 299–333. Amsterdam: John Benjamins.
- Glynn, D. Forthcoming. The semantics of extralinguistic variation. A quantitative study of dialect effects on polysemy. In *Cognitive Sociolinguistics*, G. Kristiansen, D. Geeraerts & Y. Peirsman (eds). Berlin: Mouton de Gruyter.
- Glynn, D. In press. Lexical fields, grammatical constructions, and synonymy. A study in usage-based Cognitive Semantics. In *Cognitive Foundations of Linguistic Usage-Patterns*, H.-J. Schmid & S. Handl (eds). Berlin: Mouton de Gruyter.
- Glynn, D. & Fischer, K. (eds). In press. *Usage-Based Cognitive Semantics. Corpus-Driven Approaches to Conceptual Structure*. Berlin: Mouton de Gruyter.
- Glynn, D. & Robinson, J. (eds). In press. *Polysemy and Synonymy. Corpus Methods and Applications in Cognitive Linguistics*. Amsterdam: John Benjamins.
- Gries, St. Th. 2003. *Multifactorial Analysis in Corpus Linguistics: A Study of Particle Placement*. London: Continuum.

- Gries, St. Th. 2006. Corpus-based methods and cognitive semantics: The many senses of *to run*. In *Corpora in Cognitive Linguistics. Corpus-Based Approaches to Syntax and Lexis*, St. Th. Gries & A. Stefanowitsch (eds), 57–100. Berlin: Mouton de Gruyter.
- Gries, St. Th. & Divjak, D. This Volume. Behavioral profiles: a corpus-based approach to cognitive semantic analysis. In *New Directions in Cognitive Linguistics*, V. Evans & S. Pourcel (eds). Amsterdam: John Benjamins.
- Gries, St. Th. & Stefanowitsch, A. 2004. Extending collocation analysis: A corpus-based perspective on 'alternations'. *International Journal of Corpus Linguistics* 9: 97–129.
- Gries, St. Th. & Stefanowitsch, A. (eds). 2006. *Corpora in Cognitive Linguistics. Corpus-Based Approaches to Syntax and Lexis*. Berlin: Mouton de Gruyter.
- Gries, St. Th. & Stefanowitsch, A. In Press. Cluster analysis and the identification of collocational classes. In *Empirical and Experimental Methods in Cognitive/Functional Linguistics*, J. Newman & S. Rice (eds). Stanford: CSLI.
- Grondelaers, S. & Geeraerts, D. 2003. Towards a pragmatic model of cognitive onomasiology. In *Cognitive Approaches to Lexical Semantics*, H. Cuyckens, R. Dirven & J. R. Taylor (eds), 67–92. Berlin: Mouton de Gruyter
- Grondelaers, S., Geeraerts, D. & Speelman, D. 2007. A case for cognitive corpus linguistics. In *Methods in Cognitive Linguistics*, M. Gonzalez-Marquez, I. Mittelberg & M. Spivey (eds), 149–170. Amsterdam: John Benjamins.
- Herskovits, A. 1988. Spatial expressions and the plasticity of meaning. In *Topics in Cognitive Linguistics*. B. Rudzka-Ostyn (ed.), 271–297. Amsterdam: John Benjamins.
- Heylen, K. 2005. A quantitative corpus study of German word order variation. In *Linguistic Evidence. Empirical, Theoretical and Computational Perspectives*, S. Kepser & M. Reis (eds), 241–264. Berlin: Mouton de Gruyter.
- Heylen, K., Tummars, J. & Geeraerts, D. 2008. Methodological issues in corpus-based Cognitive Linguistics. In *Cognitive Sociolinguistics. Language Variation, Cultural Models, Social Systems*, G. Kristiansen & R. Dirven (eds), 91–128. Berlin: Mouton de Gruyter.
- Hottenroth, P.-M. 1991. Präpositionen und Objektkonzepte. Ein kognitiv orientiertes, zweistufiges Modell für die Semantik lokaler Präpositionen. In *Approaches to Prepositions*, G. Rauh (ed.), 9–28. Tübingen: Narr.
- Janda, L. 2007. What makes Russian bi-aspectual verbs special. In *Cognitive Paths into the Slavic Domain*. D. Divjak & A. Kochanska (eds), 83–111. Berlin: Mouton de Gruyter
- Janda, L. & Solovyev, V. Forthcoming. Antonyms as near-synonyms: 'sadness' and 'happiness' in Russian. In *Polysemy and Synonymy. Corpus Methods and Applications in Cognitive Linguistics*, D. Glynn & J. Robinson (eds). Amsterdam: John Benjamins.
- Kemmer, S. 2006. Lexical semantics and corpora: English dimensional terms. Keynote Lecture at the First International Conference of the Swedish Cognitive Linguistics Association. Umeå, Sweden.
- Kleiber, G. 1983. Polysémie et référence: la polysémie, un phénomène pragmatique? *Cahiers de lexicologie* 44: 85–103.
- Kleiber, G. 1999. *Problèmes de sémantique: la polysémie en questions*. Villeneuve-d'Ascq: Septentrion.
- Kreitzer, A. 1997. Multiple levels of schematization: A study in the conceptualization of space. *Cognitive Linguistics* 8: 291–325.
- Lakoff, G. 1987. *Women, Fire, and Dangerous Things. What categories can reveal about the mind*. Chicago: University of Chicago Press.
- Langacker, R. 1988. A usage-based model. In *Topics in Cognitive Linguistics*, B. Rudzka-Ostyn (ed.), 127–161. Amsterdam: John Benjamins.
- Langacker, R. 2000. A dynamic usage-based model. In *Usage Based Models of Language*, S. Kemmer & M. Barlow (eds), 1–63. Stanford: CSLI.
- Langacker, R. This Volume. Constructions and constructional meaning. In *New Directions in Cognitive Linguistics*, V. Evans & S. Pourcel (eds). Amsterdam: John Benjamins.

- Lehrer, A. 1982. *Wine and Conversation*. Bloomington: Indiana University Press.
- Lehrer, A. 1990. Polysemy, conventionality, and the structure of the lexicon. *Cognitive Linguistics* 1: 207–246.
- Lemmens, M. 1998. *Lexical Perspectives on Transitivity and Ergativity. Causative Constructions in English*. Amsterdam: John Benjamins.
- Lemmens, M. 2006. Caused posture. Experiential patterns emerging from corpus research. In *Corpora in Cognitive Linguistics. Corpus-Based Approaches to Syntax and Lexis*, St. Th. Gries & A. Stefanowitsch (eds), 263–298. Berlin: Mouton de Gruyter.
- Lewandowska-Tomaszczyk, B. 1998. Dynamic perspective on antonymous polysemy. In *Making Meaningful Choices in English*, R. Schulze (ed.), 121–137. Tübingen: Narr.
- Lewandowska-Tomaszczyk, B. 1999. The 'tertium comparationis' and the concept of equivalence. In *Cognitive Perspectives on Language*, B. Lewandowska-Tomaszczyk (ed.), 53–76. Frankfurt: Lang.
- Lewandowska-Tomaszczyk, B. 2002. Polysemy: Research methodology and mechanisms. In *Cognitive Linguistics Today*, B. Lewandowska-Tomaszczyk & K. Turewicz (eds), 81–96. Frankfurt: Lang.
- Lewandowska-Tomaszczyk, B. 2007. Polysemy, prototypes, and radial categories. In *Handbook of Cognitive Linguistics*, D. Geeraerts & H. Cuyckens (eds), 139–169. Oxford: Oxford University Press.
- Lindner, S. 1983. *A Lexico-Semantic Analysis of English Verb-Particle Constructions*. LAUT: Trier.
- Meex, B. 2001. The spatial and non-spatial sense of the German preposition *über*. In *Polysemy in Cognitive Linguistics*, H. Cuyckens & B. Zawada (eds), 1–36. Amsterdam: John Benjamins.
- Newman, J. & Rice, S. 2004. Patterns of usage for English *STAY*, *STAND*, and *LIE*: A cognitive exploration in corpus linguistics. *Cognitive Linguistics* 15: 351–396.
- Newman, J. & Rice, S. 2006. Transitivity schemas of English *EAT* and *DRINK* in the BNC. *Corpora in Cognitive Linguistics. Corpus-Based Approaches to Syntax and Lexis*, St. Th. Gries & A. Stefanowitsch (eds), 225–260. Berlin: Mouton de Gruyter.
- Newman, J. & Rice, S. (eds). In press. *Empirical and Experimental Methods in Cognitive/Functional Research*. Stanford: CSLI.
- Picoche, J. 1994. A 'Continuous Definition' of polysemous items: Its basis, resources and limits. In *Continuity in Linguistic Semantics*, C. Fuchs & B. Victorri (eds), 77–107. Amsterdam: John Benjamins.
- Rice, S. 1999. Aspects of prepositions and prepositional aspect. In *Issues in Cognitive Linguistics*, L. De Stadler & Chr. Eyrych (eds), 225–247. Berlin: Mouton de Gruyter.
- Rice, S., Sandra, D. & Vanrespaille, M. 1999. Prepositional semantics and the fragile link between space and time. In *Cultural, Psychological and Typological Issues in Cognitive Linguistics*, M. K. Hiraga, C. Sinha & S. Wilcox (eds.), 107–127. Amsterdam: John Benjamins.
- Radden, G. 1981. Die übertragenen Bedeutungen der englischen Raumpräpositionen. In *Kassusgrammatik und Fremdsprachendidaktik*, G. Radden & R. Dirven (eds), 133–179. Trier: Wissenschaftlicher.
- Rémi-Giraud, S. & Panier, L. 2003. *La Polysémie ou l'Empire des Sens. Lexique, Discours, Représentations*. Lyon: Press Universitaire de Lyon.
- Rudzka-Ostyn, B. 1989. Prototypes, schemas, and cross-category correspondences: The case of *ask*. In *Prospects and Problems of Prototype Theory*, D. Geeraerts (ed.), 613–661. [Linguistics 27]. Berlin: Mouton de Gruyter.
- Rudzka-Ostyn, B. 1995. Metaphor, schema, invariance: The case of verbs of answering. In *By Word of Mouth: Metaphor, Metonymy, and Linguistic Action in a Cognitive Perspective*, L. Goossens, P. Pauwels, B. Rudzka-Ostyn, A.-M. Simon-Vandenberghe & J. Vanparys (eds), 205–244. Amsterdam: John Benjamins.
- Sandra, D. 1998. What linguists can and can't tell you about the human mind. A reply to Croft. *Cognitive Linguistics* 9: 361–378.
- Sandra, D. & Cuyckens, H. 1999. Fuzzy at first sight: The case of two Dutch prepositions. In *Cognition and Function in Language*, B. Fox, D. Jurafsky & L. Michaelis (eds), 200–219. Stanford: CSLI.
- Sandra, D. & Rice, S. 1995. Network analyses of prepositional meaning: Mirroring whose mind – the linguist's or the language user's? *Cognitive Linguistics* 6: 89–130.

- Schmid, H. J. 1993. *Cottage und Co., idea, start vs. begin: Die Kategorisierung als Grundprinzip einer differenzierten Bedeutungsbeschreibung*. Tübingen: Niemeyer.
- Schmid, H.-J. 2000. *English Abstract Nouns as Conceptual Shells. From Corpus to Cognition*. Berlin: Mouton de Gruyter.
- Schulze, R. 1988. A short story of down. In *Understanding the Lexicon. Meaning, Sense, and World Knowledge in Lexical Semantics*, W. Hüllen & R. Schulze (eds), 395–414. Tübingen: Niemeyer.
- Schulze, R. 1994. Image schemata and the semantics of *off*. In *Kognitive Semantik. Ergebnisse, Probleme, Perspektiven*, M. Schwartz (ed.), 197–213. Tübingen: Narr.
- Stefanowitsch, A. & Gries, St. Th. 2003. Collostructions: Investigating the interaction between words and constructions. *International Journal of Corpus Linguistics* 8: 209–243.
- Stefanowitsch, A. & Stefan Th. Gries (eds). 2006. *Corpus-based Approaches to Metaphor and Metonymy*. Berlin: Mouton de Gruyter.
- Talmy, L. 2000. *Toward a Cognitive Semantics*. Cambridge MA: MIT Press.
- Taylor, J. 1993. Prepositions: Patterns of polysemisation and strategies of disambiguation. In *The Semantics of Prepositions*, C. Zelinsky-Wibbelt (ed.), 151–175. Berlin: Mouton de Gruyter.
- Taylor, J. 2003a. Cognitive models of polysemy. In *Polysemy – Flexible Patterns of Meaning in Mind and Language*, B. Nerlich, Z. Todd, V. Herman & D. D. Clarke (eds), 31–48. Berlin: Mouton de Gruyter.
- Taylor, J. 2003b. Polysemy's paradoxes. *Language Sciences* 25: 637–655.
- Tuggy, D. 1993. Ambiguity, polysemy, vagueness. *Cognitive Linguistics* 4: 148–162.
- Tuggy, D. 1999. Linguistic evidence for polysemy in the mind. *Cognitive Linguistics* 10: 343–368.
- Tummers, J., Heylen, K. & Geeraerts, D. 2005. Usage-based approaches in Cognitive Linguistics. A technical state of the art. *Corpus Linguistics and Linguistic Theory* 1: 225–261.
- Tyler, A. & Evans, V. 2001. Reconsidering prepositional polysemy networks: The case of *over*. *Language* 77: 724–65.
- Tyler, A. & Evans, V. 2003. *The Semantics of English Prepositions: Spatial Scenes, Embodied Meaning, and Cognition*. Cambridge: Cambridge University Press.
- Vandeloise, C. 1984. *L'Espace en Français*. Paris: Seuil.
- Vandeloise, C. 1990. Representation, prototypes, and centrality. In *Meanings and Prototypes. Studies on Linguistic Categorization*, S. Tsohatzidis (ed.), 403–437. London: Routledge.
- Victorri, B. & Fuchs., C. 1996. *La Polysémie: Construction Dynamique du Sens*. Paris: Hermès.
- Wulff, S. 2006. Go-V vs. go-and-V in English: A case of constructional synonymy? In *Corpora in Cognitive Linguistics. Corpus-Based Approaches to Syntax and Lexis*, St. Th. Gries & A. Stefanowitsch (eds), 101–126. Berlin: Mouton de Gruyter.
- Wulff, S., Stefanowitsch, A. & Gries, St. Th. 2007. Brutal Brits and persuasive Americans: variety-specific meaning construction in the *into*-causative. In *Aspects of Meaning Construction*, G. Radden, K.-M. Köpcke, T. Berg & P. Siemund (eds.), 265–281. Amsterdam: John Benjamins.
- Zelinsky-Wibbelt, C. 1986. An empirically based approach towards a system of semantic features. *Proceedings of the 11th International Conference on Computational Linguistics* 11: 7–12.
- Zelinsky-Wibbelt, C. 1993. Interpreting and translating prepositions: A cognitively based formalization. In *The Semantics of Prepositions: From Mental Processing to Natural Language Processing*, C. Zelinsky-Wibbelt (ed.), 351–390. Berlin: Mouton de Gruyter.
- Zelinsky-Wibbelt, C. 2000. *Discourse and the Continuity of Reference. Representing Mental Categorization*. Berlin: Mouton de Gruyter.
- Zeschel, A. (ed.). 2008. *Usage-Based Approaches to Language Processing and Representation*. Special Edition of *Cognitive Linguistics* 19.
- Zlatev, J. 2003. Polysemy or generality? Mu. In *Cognitive Approaches to Lexical Semantics*, H. Cuyckens, R. Dirven & J. R. Taylor (eds), 447–494. Berlin: Mouton de Gruyter.

PART II

Approaches to metaphor and blending

Theory and method

Solving the riddle of metaphor*

A salience-based model for metaphorical interpretation in a discourse context

Mimi Ziwei Huang

At dusk I come without being fetched. At dawn I disappear without being stolen. I am a poet's tears and a sailor's guide. What am I? (An English folk riddle)

1. What do we need to interpret a metaphorical riddle?

The English folk riddle quoted at the beginning of this paper is not difficult, and most of us can guess the answer to be “the stars”. While stars can literally be “a sailor’s guide”, they are only “a poet’s tears” when understood metaphorically. Only when we interpret stars figuratively as people can the first person point of view “I” make sense in the riddle. When guessing the meaning of this riddle, we also undergo a cognitive process of interpreting metaphors. To explain such a process of comprehension is to reveal our cognitive principles in constructing a metaphorical meaning. The aim of this paper is to explore how metaphorical meaning is developed and communicated in a discourse.

Recent studies in cognitive linguistics have offered differing explanations for the meaning construction and cognitive processes involved in metaphor (e.g., Lakoff and Johnson 1980; Fauconnier and Turner 2002; Giora 2003; Carston 2002; Evans Forthcoming), and have answered many important questions in metaphorical comprehension. It is not the aim of this paper to address all the aspects of metaphor. Rather, it will tackle one particular issue: the construction of metaphorical meaning in a given discourse. More specifically, I will focus on a salience-based model, which builds upon a notion of salience defined by degrees of accessibility. This salience-based model, as I will argue in the following, provides a useful analytical approach for examining the development of a metaphorical meaning.

This notion of salience will be examined in Section 2. I will review in this section an established salience framework: the Graded Salience Hypothesis and its two supplementary hypotheses (Giora 1997, 1999; Peleg et al. 2001). Giora’s approach to metaphorical

* I would like to thank Rachel Giora for her valuable comments on the NDCL presentation which this paper is based on.

meanings will be examined through examples, which will lead to a further discussion of the notion of salience in metaphorical analysis. Based on this review, Section 2 will propose a modified notion of salience, which is coherent with its counterpart in the Graded Salience Hypothesis, yet also differs substantially from it. This proposed notion of salience will be elaborated upon in Section 3 where it will be discussed from a variety of cognitive aspects. This proposed notion of salience will be employed in Section 4 to analyze the metaphorical interpretation of a short story. Finally Section 5 will summarize the discussions in this paper.

2. The notion of salience and its role in metaphorical interpretations

The notion of salience in this paper is closely related to the Graded Salience Hypothesis (Giora 1997, 2003). Giora and her colleagues have been interested in meaning activation and construction in discourse processing, and have developed the Graded Salience Hypothesis to explore if meaning is accessed in a hierarchical manner in cognitive processing, including metaphorical comprehension. The contention of the Graded Salience Hypothesis is that salient – i.e. consolidated and encoded – lexical meanings of a mental entity are always activated in the initial process of comprehension, regardless of the context. The Graded Salience Hypothesis has gained support from various empirical studies, some of which include reading times, response times, moving windows, eye tracking and scanning brain waves (see Katz and Ferretti 2000, 2001; Anaki et al. 1998; Pexman et al. 2000; Turner and Katz 1997). The Graded Salience Hypothesis is further supplemented by the Retention Hypothesis (Giora and Fein 1999) and the Optimal Innovation Hypothesis (Giora et al. 2004), which will also be reviewed in the following.

2.1 The graded salient lexical meanings

According to the Graded Salience Hypothesis, salience is defined according to its consolidation in the mental lexicon. Salient meanings are the encoded lexical meanings of a word or an expression that are high in usage frequency, familiarity, conventionality, and prototypicality/stereotypicality (Giora 1997, 2003; Giora and Fein 1999). Only consolidated and lexicalised meanings are salient, and meanings that are not coded in the mental lexicon are *nonsalient* (Giora 2002: 490, italics as in the original).

The Graded Salience Hypothesis regards lexical saliency as a matter of degree rather than an absolute attribute of a word or an expression (Giora 2003: 15). Also, a word or an expression can have more than one salient lexical meaning, and if these meanings are similarly salient, they will all be accessed simultaneously (Giora 2003: 37).

In metaphorical comprehension, the Graded Salience Hypothesis predicts that literal meanings will always be activated in different types of context due to their high saliency. Metaphorical meanings in familiar metaphors are also encoded, and enjoy a similar degree of salience to literal meanings, and will always be activated as well. An utterance such as ‘I’ve got cold feet’, with the expression ‘cold feet’ being a familiar metaphor, will trigger

the metaphorical meaning of “fearfulness or timidity” even when the intended meaning is the literal one (Giora 2003). In novel metaphors, however, because the metaphorical meanings are not encoded in the mental lexicon, they will only be activated in metaphorically-biasing contexts, but not in literally-biasing ones.

2.2 Salient lexical meanings and context

The stage where salient lexical meanings are first activated is the initial process, which is followed by the integration process (Giora 1997, 2003; Peleg et al. 2001). At the initial stage, contextual information is processed in parallel with the mental lexicon, which will not be inhibited even if it is contextually incompatible.

In the second stage, the integration process takes place, and is governed by the Retention Hypothesis, which predicts that the initially activated salient meanings will be sustained unless they are intrusive to or conflict with the contextual information (see Giora and Fein 1999 for discussion of the Retention Hypothesis). At this stage, the Graded Salience Hypothesis is also complemented by the Optimal Innovation Hypothesis, which allows “for the recovery of a salient meaning from which that novel meaning stems, in order that the similarity and difference between them may be assessable” (Giora 2003: 176; see also Giora et al. 2004). The Innovation Hypothesis predicts the retrieval of the original salient meanings which may be useful or necessary in meaning processing (see also Sperber and Wilson 1995; Evans Forthcoming for discussions of the role of literal meanings in constructing metaphorical interpretation).

2.3 A discussion of Giora’s salience framework

The Graded Salience Hypothesis has made a valuable contribution to the study of metaphor. While I concur to a large extent with Giora’s discussion of the Graded Salience Hypothesis and the two supplementary hypotheses (Giora 2003), I wish to further examine the notion of salience in terms of accessibility of meaning in a given discourse. I will in this section address two issues in Giora’s salience framework, the first of which concerns Giora’s examination of metaphorical meaning, and the second relates to Giora’s discussion of literal meaning of a metaphor.

As commonly defined, a metaphor depicts a target entity in terms of a source entity. In most circumstances, the metaphorical meaning would to some extent modify the meaning of the target entity as the result of the metaphorical mapping process. The degrees of modification may vary, but the cognitive connection between the literal meaning and the metaphorical meaning can be traced. In Giora’s discussion of metaphors, the source entities are mostly thoroughly examined and tested for the activation of their salient meanings. Meanings of the target entities, however, are not always examined to a similar extent. Semino (2004: 2188), for instance, has suggested that Giora’s approach to the meaning of words and expressions could be problematic. Semino quotes one example as illustrated in Giora’s discussion of metaphor (Giora 2003: 112, italics and asterisk as in the original):

- (1) Sarit's sons and mine went on fighting continuously. Sarit said to me: These *delinquents** won't let us have a moment of peace.
 (Probes displayed at *: Salient: criminals; Contextually compatible: kids; Unrelated: painters; Nonword: nimvhar)

Example (1) above is treated as a novel metaphor. Using the single-word probe method, Giora investigates the “priming effects” after the target word “delinquents”. According to Giora, the word “delinquents” in Example (1) would have two meanings: its literal and salient meaning “criminals”, and its metaphorical and non-salient meaning “kids”. The former is the unintended meaning, and hence contextually incompatible, the latter is the intended meaning, and hence contextually compatible. Participants’ response times for the above four probes are recorded for comparison. The result shows that responses to the salient but contextually incompatible probes (criminals) do not differ from the non-salient but contextually compatible probes (kids). This result supports the prediction of the Graded Salience Hypothesis in the sense that the literal meaning of a novel metaphor will be activated initially, regardless of the context.

Giora’s analysis is not entirely agreed upon by Semino (2004), who questions the metaphorical meaning in Example (1). Semino notes that the metaphorical meaning “kids” as suggested by Giora is not adequate when interpreting the overall meaning of Example (1). According to Semino, at least some aspects of the salient lexical meaning of “delinquents” are applied to the children in question, and via a cross-domain mapping, achieve a complete understanding of the metaphor, which does not simply mean “kids”, but rather, “our excessively badly behaved, aggressive, litigious, etc., kids”. Semino argues that it is only in a very strict sense that the literal meaning of “delinquents” can be described as contextually incompatible (2004: 2188). Giora’s discussion of Example (1) has revealed a cognitive connection between the source entity *DELINQUENTS* and the target entity *KIDS*,¹ but it does not illustrate how the salient meaning of the entity *DELINQUENTS* helps modify the entity *KIDS* in an utterance such as Example (1).

Semino’s observation raises an important question in the study of metaphor: How is the metaphorical meaning constructed? When a salience-based account is applied to analyze metaphor, there comes a related question: How accessible is the metaphorical meaning in a given discourse? The first question – the construction of the metaphorical meaning – does not seem to be central in Giora’s discussion of metaphor. The metaphorical meaning is very often available already in Giora’s discussion for the analytical purpose (such as the suggested metaphorical meaning in the above example). The second question – the accessibility of the metaphorical meaning – is discussed in terms of salience in Giora’s approach. The metaphorical meaning, when being the intended meaning, should be accessible to the reader, but is only regarded as salient if encoded in the reader’s mental lexicon.

Now let us examine Example (1) from a different aspect, and see if the above two questions can be answered from a different aspect. In (1), the metaphorical meaning “kids” as suggested by Giora is not an encoded meaning of the entity *DELINQUENTS*. Therefore although it is an accessible meaning in (1), it is not salient in Giora’s framework. However,

1. Small capitals are used in this paper to refer to mental entities.

before the reader processes (1), the metaphorical meaning “kids” is not available in the comprehension of DELINQUENTS. In other words, this meaning is not accessible. After (1) is processed, the meaning “kids” emerges as a newly accessible meaning. Although it is not consolidated (yet) in the reader’s cognition, its degree of salience has undergone a change, from being not-accessible to accessible, and perhaps even highly accessible when restricted to the local text such as (1). A newly-accessible meaning as such can be retained for further processing should the contextual information support it. Suppose following Example (1), after Sari said to me “These *delinquents* won’t let us have a moment of peace”, I then said “You are right. I think we should teach the delinquents a lesson.” The metaphorical meaning “delinquent-like kids” is retained in the given discourse. Now suppose this metaphorical meaning becomes a regular use to the speakers, and it gets to a stage when the word “delinquent” is uttered, the speakers will naturally think of their kids fondly as well. (Let us hope one would not encounter real delinquents all the time in real life!) It can be said at this stage that the metaphorical meaning of “delinquent” becomes also highly accessible at a personal level. Now let us take one step further and suppose the metaphorical usage of “delinquents” has become widely acceptable and has gone through a consolidation process, the meaning of “delinquent-like kids” will then gain a high accessibility at a global level. At this stage, such a meaning would also be regarded as “salient” according to the Graded Salience Hypothesis.

This alternative approach to examining the graded accessibility of a modified meaning can be a useful supplement to Giora’s salience-based framework. The main difference between this alternative approach and Giora’s framework is that the latter employs consolidation and encodedness in meaning construction as a benchmark in the discussion of salience while the former does not. Before further comparison between these two approaches, I will, as mentioned at the beginning of this section, examine another issue that relates to salience – Giora’s approach to the literal meaning of a metaphor.

In Giora and Balaban (2001), written texts from newspaper articles that contain metaphors are selected to test the notion of salience in the Graded Salience Hypothesis, which predicts that the literal meanings of familiar metaphors, albeit being contextually incompatible, is lexically salient and will not be blocked by the context. Among the 60 metaphors selected, 30 of them have some mention or echo of their unintended literal meaning, which is represented by a word semantically related to their literal meanings, such as in Examples (2) and (3) below (Giora and Balaban, 2001: 117, 120, italics and ellipses as in the original); the other half of the selected metaphors did not have their unintended literal mentioned or evoked, such as in Example (4) (*ibid*, 117):

- (2) The strikes in the education system took place when the union was putting up a *fight* against the government. In this *fight*, threats, sanctions, and even a general strike were the *weapons*.
- (3) Israel needs ... not only those who *flirt* with the capital market, but those who *marry* it, for *better or worse, in poverty and in wealth, until “a purchase proposal” do they apart*.
- (4) He *lost* his health, and his spirit *broke*.

These texts are read by a total of 40 participants who rate the metaphors on a 1–7 familiarity scale. The results, according to Giora and Balaban's interpretation, are consistent with the predictions made by the Graded Salience Hypothesis. Metaphors, such as the use of "fight" in (2) and "flirt" in (3), are each followed by the mention of their literal meanings, i.e. "weapons" in (2) and "marry" in (3). These metaphors do not differ in familiarity from those that were not followed by the literal meaning, i.e. "lost" and "broke" in (4). These results are used to support the Graded Salience Hypothesis in the sense that the context does not block activation of salient though incompatible meanings even in highly conventional metaphors (Giora and Balaban 2001; Giora 2003: 136).

My hesitation in agreeing with Giora and Balaban lies in their examination of the "literal extension" of a metaphorically-used word. Take Example (2) for instance, the use of "fight" is a conventional metaphor, and so is the use of "weapons". The words "fight" and "weapons" are not only semantically related in their literal sense, but are also metaphorically connected in the conceptual metaphor ARGUMENT IS WAR (Lakoff and Johnson 1980). In Giora and Balaban's study (2001), Example (2) is rated by the readers as conventional, but the rating does not necessarily reveal that the readers deliberately process the literal meaning of "fight" or "weapons". What the results show is the readers do not seem to employ extra processing effort in interpreting two metaphorical concepts that both realise the same conceptual metaphor. In a similar vein, the familiar usage of "flirt" and "marry" in Example (3) both realise the MARRIAGE metaphor often found in political discourse (see Musolf, 2004 for the discussion of metaphor scenarios employed in political discourse). The readers' familiar rating of Example (3) does not directly indicate any active functions of the literal meaning of "flirt" or "marry" in the interpretation of (3). However, this is not to say that the readers always do not activate the literal meanings in the above examples. The literal meanings, being encoded and consolidated, are easily accessible should the cognitive process requires it. Also, cognitive process may vary at personal, cultural and discursal levels, and consequentially a various degrees of accessibility would be applied to even the encoded literal meaning of a conventional metaphor.

The purpose of the above discussion is not to take issue with the Graded Salience Hypothesis and its related hypotheses, but it leads to the suggestion that in metaphorical analysis, the notion of salience can be measured by an alternative indicator – dynamic accessibility – instead of meaning-consolidation in the mental lexicon. With this alternative indicator, the most salient information is the most accessible one at a given time point of cognitive processing. The accessibility is dynamic in the sense that the accessibility of a meaning varies at different cognitive levels. Some information is globally salient and its degree of accessibility is reasonably stable. Some information is only locally salient and its accessibility can undergo dramatic change in a local discourse. Some information can be personally salient and is only accessible to certain individuals but not to others. This alternative view of salience shares a number of features with the notion of salience proposed in the Graded Salience Hypothesis (Giora 1997, 2003), but also stands alone as an independent concept. Section 3 will elaborate upon this proposed notion of salience.

3. Proposal of a salience-based framework for metaphorical analysis

3.1 The alternative notion of salience

The previous section has provided an initial discussion of a salience-based framework that features a notion of salience that differs from the one proposed in Giora's framework (Giora 1997, 1999, 2003). In this paper and in Huang (2008), the meaning of a mental entity is considered salient if it is highly accessible at a given time point in a given discourse. This can be used as a working definition of the alternative notion of salience. From this section onwards, unless otherwise specified, "salience" is referred to in the above defined sense. The notion of salience in Giora's framework is specified as "encoded salience" due to its employment of meaning-consolidation as the definitive benchmark.

Salience is a dynamic and graded notion. It varies at differing stages of processing and in differing discourses. It also reflects the situation where one mental entity has multiple meanings competing for salience at the same time. Importantly, salience functions in a discourse, which is composed of two core elements: the formal linguistic quality of a particular stretch of language, and its individual or group users. (see Cook 1994: 25–27 for the discussion of discourse). In this sense, salient meanings are subjected to the range of linguistic and extra linguistic information and knowledge accessible to the language users.

The main difference between this notion of salience and the encoded salience (Giora 2003) is that while both notions regard encoded meaning to be salient, the former also regards non-consolidated meaning to be salient at a given time point if it is highly accessible. This notion of salience provides the basis for a salience-based model which examines metaphorical interpretation in a discourse context. A detailed discussion of this framework is presented in Huang (2008). Due to the length of this paper, I will only discuss some of the important aspects of the framework that is particularly relevant to the discussion in this paper.

Similar to the Graded Salience Hypothesis (Giora 2003), the salience-based model proposed in Huang (2008) regards encoded meaning of a mental entity to be salient. Encoded meanings are easily retrievable in one's mental lexicon. When a mental entity is being processed, even if its encoded meaning is not the most accessible at a given time point, it can still be easily activated should it be required. I will in Section 4 explain the method for identifying the encoded meaning of a mental entity. In this section, I will focus on the examination of information that is non-consolidated yet salient in cognitive process, and discuss the importance of this information in metaphorical interpretation.

In Huang (2008), four types of non-consolidated information are discussed with regards to their potential salience in cognitive process. Differing from encoded information that can be directly retrieved from one's mental lexicon regardless of context (Giora 2003), non-consolidated meaning cannot be directly accessed or retrieved in one's cognitive process. However, non-consolidated meaning can be *made* highly accessible in a given discourse, and they can include four types of information: (1) the meaning with syntactic prominence in the given text; (2) given a meaning already being salient in a discourse, its

semantically and/or schematically most related meaning that is previously mentioned in the same discourse; (3) the foregrounded meaning in the given discourse; (4) the most preferable and/or desirable meaning considered by the discourse participant. I will discuss each of these meanings in the following.

3.2 Syntactically salient meaning

The consideration of syntactic salience in the salience-based model for metaphor (Huang 2008, Ref. Chapter 4) is inspired by recent studies in the field of Information Structure (Grosz et al. 1995; Walker et al. 1998; and Hajičová et al. 1998b; to name a few). In a nutshell, in the syntactic aspect, the salience of a mental entity depends on the position of its linguistic representation in the sentence. In an utterance *X is Y*, the meaning of *Y* is more accessible than the meaning of *X* because *Y* is the syntactic focus and *X* is the syntactic topic. The syntactic topic often carries the given information and the syntactic focus carries the new information, and therefore the focus has new information to “talk about” the topic (see Hajičová et al. 1998 for the discussion of syntactic topic and focus).

The syntactic form “*X is Y*” is also a prototypical linguistic expression for a metaphor, in which *X* stands for the target entity and *Y* stands for the source entity. Take Example (5) for instance:

- (5) Life is a box of chocolates.

Being the syntactic focus, the meaning of “a box of chocolates” is more accessible than the meaning of “life” when (5) is being processed. The metaphorical meaning “a box of chocolates” is probably not consolidated in one’s mental lexicon; even if it is, the meaning is not more encoded than the literal meaning of “life”. However, the metaphorical meaning of “a box of chocolates” is, at the local textual level, situated in the centre of the reader’s attention, and is hence more salient than the meaning of “life”. Syntactic salience as such is a local notion. Strictly speaking, a meaning is syntactically salient only when the particular sentence is processed. Albeit being a local notion, syntactic salience is a useful indicator in monitoring the degree of salience of a mental entity in a given discourse. Section 4 will outline a topic-focus analytical framework that traces the development of a salient meaning in metaphorical interpretation.

3.3 Salience in associate entities

Apart from syntactic salience, the meaning of a salient entity’s immediately associate entity can also gain a reasonably high degree of salience regardless of its status of consolidation in one’s mental lexicon.

Salience in associate entities supplements syntactic salience discussed previously. In a given discourse, if an immediate association exists between two entities (i.e. that there is a sufficient body of relevance assembled between them), these two entities are associate entities. (See Hajičová and Sgall 2001; Hajičová et al. 2003 for discussion of associate entities.) Generally speaking, two words that belong to the same semantic fields are likely

to represent two entities that are associated with each other. Hajičová et al. (2003) give the example that the entity CHILD is immediately associated to the entity PARENTS OR MOTHER. Similarly, Example (2) quoted in Section 2.3 of this paper can be said to realise the entities FIGHT and WEAPONS, which are also associate entities when Example (2) is processed.

Hajičová et al. (Hajičová and Sgall 2001; Hajičová et al. 2003) argue that in a given text (particularly a written text), if an entity A is syntactically salient when sentence S is processed, and entity B is associated with A and is mentioned in a previous sentence, the salience degree of B at the point when S is processed would be two degrees lower than that of A (see Hajičová et al. 1998a: 383–384 for a tentative rules for assigning salience degrees to mental entities in a given discourse). Take Example (2) again for instance, if the metaphorical meaning of “weapons” is the most salient information when (2) is read, then the metaphorical meaning of “fight”, which is activated in the previous sentence, will be less salient than that of “weapons”, but still highly accessible. In a given discourse, if a metaphorical meaning sustains its salience, it would indicate a continuous metaphorical interpretation in the reader’s cognitive process. This matter will be further discussed in Section 4 in the following.

3.4 Salience in the foregrounding meaning

So far I have examined two types of information that do not depend on meaning consolidation. The third type of non-encoded meaning that can become salient in a discourse is foregrounded meaning which develops out of the reading process.

The notion of foregrounding discussed in van Peer and Hakemulder (2006) applies to a wide range of cognitive aspects, which include one of the following: certain special prominence emerging from the reading act; the specific stylistic devices in the text employed by the author; the specific poetic effect on the reader; an analytic category for literary appreciation; and finally, a distinction between literary and non-literary discourse (van Peer and Hakemulder 2006: 546–547).

In Huang (2008), the above senses of foregrounding have been discussed in relation to the notion of salience. The information that is made prominent, i.e. foregrounded, in a given discourse is also the salient information discussed in this paper. In a metaphorical interpretation, particularly in the case of a novel metaphor, the newly emerged metaphorical meaning can be regarded as a type of foregrounded information that contrasts with the known information. This foregrounded meaning might not be encoded in one’s mental lexicon, but it is highly prominent, and hence salient. In Example (1) discussed previously in Section 2.3, the expression “[t]hese delinquents won’t let us have a moment of peace” is used to refer to the speaker’s children. Although the newly emerged metaphorical meaning of “delinquents” is not an encoded meaning, it is a foregrounded information, and is salient in Example (1).

3.5 Meanings preferable to the readers

The fourth type of salient information that does not depend on the entity's consolidated lexical meaning is the most preferable and/or desirable meaning considered by the discourse participants. Recent studies in psycholinguistics have suggested that readers' preferences can affect their interpretive outcomes of literary narratives (Gerrig and Rapp 2004; Rapp and Gerrig 2002; Rapp et al. 2001; Prentice et al. 1997). When several possible meanings are equally accessible in a given discourse, the readers can choose to agree or reject certain meaning according to their personal preferences.

Research findings in Huang (2008) have supported this argument. In Huang (2008, see Chapters 7, 8 and 9), three short stories with possible metaphorical interpretations are examined. The analysis of each story's textual information highlights a number of meanings that are salient because they are either the encode meanings of the mental entities represented in the discourse, and/or they are syntactically salient in the given sentences; and/or they represent associate entities of salient entities in the discourse; and/or they are the foregrounded meanings in the story. These meanings are further examined in a series of case studies where a number of readers are interviewed for their interpretations of each story. The interview data show that although the above textually salient meanings are mostly accessible to the readers in their reading process, the readers tend to employ personal preferences in their interpretations to sustain the favourable meanings, and to resist or reject the undesirable ones.

In the reading process, if one meaning is the most preferable one according to a reader, it is also the most salient meaning at the time when the reader decides her preference. This type of preferable meaning is the fourth type of non-encoded yet salient information in the salience-based model for metaphor proposed in this paper and in Huang (2008).

3.6 Summary

The above sections have outlined and discussed a notion of salience that is defined by a dynamic and multi-facet accessibility. A meaning is regarded as salient if it is the most accessible information at the given time point of discourse processing. If a meaning is encoded in a mental entity, it is salient to the reader at a global level in the sense that the reader should be able to access it in different discourses even with little aid from the contextual information. At a local textual level, the meaning with a syntactic prominence, such as the meaning conveyed by the syntactic focus, is salient at least within the range of the sentence. For a meaning that is immediately associated with another meaning that is already salient in a discourse, the associated meaning can also promote its degree of salience. In metaphorical interpretation, the metaphorical meaning is foregrounded to the reader's centre of attention, and is also salient when such a meaning becomes available in cognitive processing. Finally, in the local personal level, the meaning in a discourse that is mostly preferred by the reader is the most accessible meaning to the individual, and is personally salient to the reader when the discourse is processed.

This proposed notion of salience can be related to the notion of salience discussed in the Graded Salience Hypothesis and its two supplementary hypotheses (Giora 2003). Both notions of salience are graded, dynamic, and depend on accessibility. Encoded meanings are salient information according to both salience accounts. Non-encoded meanings are not salient according to Giora, but this locally and personally accessible information has the potential to become globally accessible should be used frequently and becomes widely recognised. A notion of dynamic salience as such can provide a useful insight to the explanation of the development of metaphorical meaning in a discourse context. In the next section, I will illustrate how the proposed notion of salience can contribute to metaphorical analysis.

4. Salience dynamics in the interpretation of a short story

In this section a short story from *The Devil's Larder* (Crace 2002) is selected for analysis. The book contains 64 individual texts numbered from 1 to 64 without specific title. Text 1 is selected for illustrative purpose here based on the grounds that the first text of the book bears limited contextual background from the previous text. Hence in the interpretive process, the reader will mainly rely on the text itself for interpretive information.

In Huang (2008), a thorough investigation of the same text has been presented. Due to the length of this paper, the examination of the text will not contain the same thoroughness. However, It is hoped that the analysis of the selected text can provide a clear illustration of how the proposed notion of salience can contribute to the study of metaphor. In the following, I will examine the selected text in terms of the five types of salient information as discussed in Section 3 above.

4.1 The encoded meaning of the CAN

For ease of reference, the following discussion will address Text 1 as *The Can Story* because the mental entity CAN is the discourse topic throughout the story. The story consists of 414 words arranged into 51 sentences and 6 paragraphs. The story is attached in the appendix for the purpose of reference. In *The Can Story*, the main characters “they” try to guess the possible contents of an unlabelled can in their larder. Starting to guess it to be ordinary food such as baked beans or fish, “they” carry on to imagine impossible contents within the can, such as “the elixir of youth” and “a devil or a god”. The entity CAN starts with its concrete, literal sense in the story, and gets more abstract towards the end. The final paragraph, Paragraph 6 (P6) quoted below with my added subscripts of mental entities represented in the paragraph,² strongly suggests a metaphorical interpretation of the entity CAN:

2. In this paper and in Huang (2008, Ref. Chapter 4), mental entities for examination are mostly represented linguistically by nouns and pronouns (see also Hajičová et al. 1998).

- (P6) We_{WE} all should have a can_{CN18} like this_{CN}. Let it_{CN20} rust. Let the rims_{RM} turn rough and brown. Lift it_{CN} up and shake it_{CN} if you want. Shake its sweetness_{SW} or its bitterness_{BT}. Agitate the juicy heaviness_{HV} within. The gravy heaviness_{HV}. The brine_{BR}, the soup_{SP}, the oil_{OL}, the sauce_{SC}. The heaviness_{HV}. The choice_{CH} is wounding it_{CN} with knives_{KF}, or never touching it_{CN} again.

According to the proposed notion of salience, the encoded meaning of a mental entity is globally salient, and is easily accessible in a given discourse. If the entity *CAN* already has an encoded metaphorical meaning, then it perhaps needs no modification in *The Can Story*. However, if such a metaphorical meaning is not yet available to the reader before the story is read, it can be argued that the metaphorical meaning, should it becomes available, is developed within the text.

According to the dictionary resources,³ the noun form “can” has twelve contemporary meanings, amongst which the meanings “airtight sealed metal container for food or drink” and “a vessel for holding liquids” are the most frequent. Both these two meanings are literal, with the core sense of being a container (Huang 2008: 230–231). There is a metaphorical meaning of “can” as referred to in the expression “can of worms”. However, the metaphorical meaning “source of unpredictable trouble and complexity” only seems to appear when the expression “can of worms” is used as a set phrase, and the frequency of this metaphorical meaning is far lower than that of the two literal meanings when the word “can” is in use (ibid.) It can be argued, therefore, when the word “can” is used individually, the most salient lexical meaning appears to be “a metal container for food, drink, paint, etc.”

When the word “can” is used in the story, this literal meaning, being the most encoded in the reader’s mental lexicon, would be most likely to be activated, and is the most salient meaning when “can” is processed at the beginning of the story.

4.2 Syntactic salience in *The Can Story*

As discussed in Section 3.2, syntactic prominence promotes the information situated in the syntactic focus of a given sentence. Also, the continuous appearance of a syntactic topic also indicates a sustaining degree of salience. Huang (2008, Chapter 4) has adopted a tentative framework to measure syntactic salience (see Hajičová et al. 1998b for the framework of the Topic-Focus Articulation).

The essence of this framework illustrates the information of a given text in a hierarchical manner. The textual information is arranged in a descending order of paragraph-entity-sentence. Sentences with the same syntactic topic are grouped under the entity that represents the topic, such as *CAN1*, *CAN2*, and *CAN3*. Entities in the focus part of the sentence will be grouped under the entity (or entities) in the topic part. Entities with multiple appearances in the text are numbered, starting with “1”. The latest mentioned entity automatically includes its previous contents, such as *CAN3* [*CAN2* (*CAN1*)] (see Huang 2008, Chapters 4 and 7 for detailed discussions of this measuring framework).

3. WordNet® and Oxford English Dictionary are used as the dictionary resources. See Huang (2008, Chapters 6 and 7) for the discussion of this methodology.

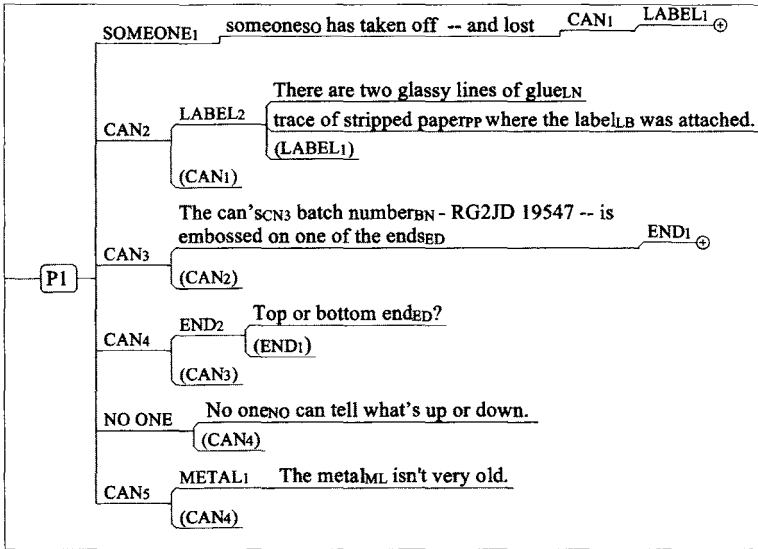


Figure 1. Textual information of *The Can Story*: Paragraph 1

With this framework, the textual information of *The Can Story* can be illustrated in Figures 1 to 5. The entities in the left columns of these figures stand for the syntactic topics of the sentences. Paragraphs 1 and 2 (P1 and P2), for instance, highlight three entities CAN, CONTENTS and THEY. These entities remain active from Paragraphs 3 to 6, which means their syntactic prominences are relatively high throughout the text. The entity THEY is replaced by the entity WE in the last paragraph (P6), which dominates most of the paragraph till the end of the story. This shift of syntactic topics indicate the narrative centre moves from the characters “they” to “we” when the story unfolds.

Importantly, Figures 1 to 5 show an intimate interweavement between the entities CAN and the CONTENTS of the can. The syntactic structure has shown that both CAN and CONTENTS are continuously modified with new information from Paragraph 1 to Paragraph 6, starting from something concrete (e.g. “salmon”, “tuna steaks” and “pineapples” in Paragraph 2) to something abstract (e.g. “plague”, “elixir of youth” and “devil” in Paragraph 4).

From the perspective of syntactic salience, when the story marches towards the end, the most salient meaning of the entity CAN is not its encoded literal meaning, but the newly emerged, abstract meaning of a metaphorical container that holds intangible contents. The syntactic information as illustrated in Figures 1 to 5 also presents an immediate cognitive association between the entities CAN and CONTENTS (of the can). The development of this pair of associate entities also helps construct the foregrounded meaning in the story. I will in the next section discuss these associate entities and the metaphorical meaning foregrounded in the story.

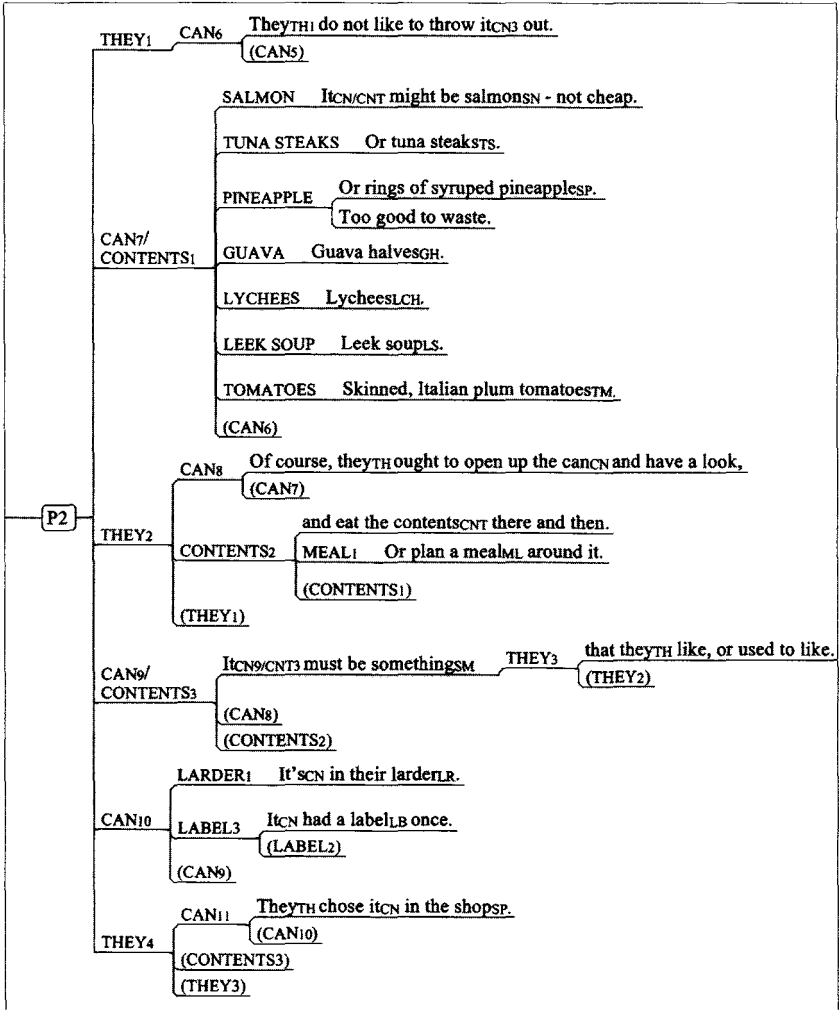


Figure 2. Textual information of *The Can Story*: Paragraph 2

4.3 Associate entities and foregrounded meaning in *The Can Story*

The cognitive association between the entity CAN and its CONTENTS is not restricted to *The Can Story*, but is instead a well-accepted notion. The association represents a metonymic relationship in which the container stands for the contents. The syntactic information in the text confirms such an immediate association. This associative relationship not only links CAN with its CONTENTS, but also indicates *what* types of CONTENTS that are normally associated with CAN. Huang (2008: 232–235) has investigated collocates and concordance of the word “can” as presented in the British National Corpus. The results show that the mostly likely contents of a “can” are indeed food-related objects such as “beans”, “coke” and “tomatoes”.

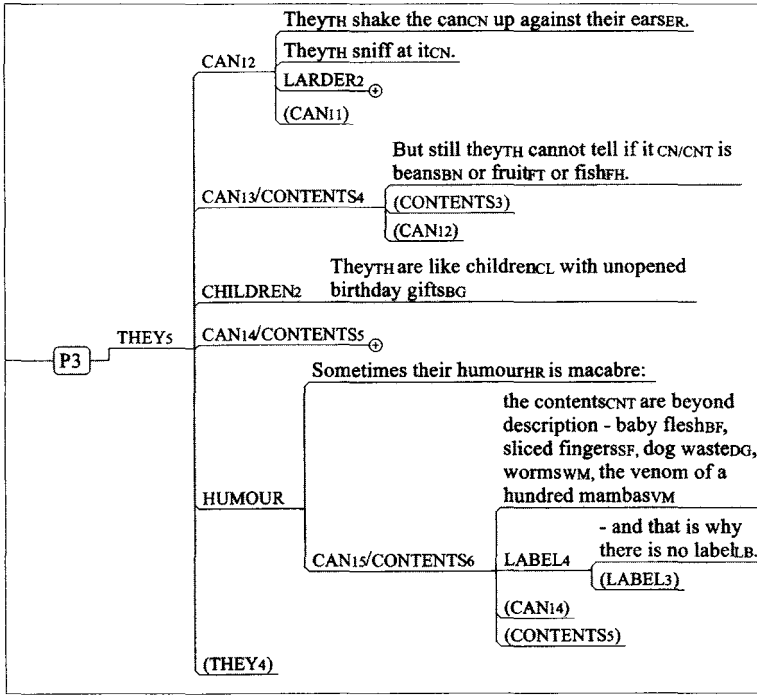


Figure 3. Textual information of *The Can Story*: Paragraph 3

This default connection between CAN and its CONTENTS is confirmed in Paragraphs 1 and 2 in *The Can Story*. Paragraphs 3 to 6, however, add unusual contents to the can, and by doing so, alters the readers' expectation of the default connection between CAN and its associate entity CONTENTS. The modification of the can's contents foregrounds an unusual situation in which only when the CAN entity is perceived metaphorically, can it accommodate those abstract and intangible contents as suggested between Paragraphs 3 to 6. In other words, only if the reader develops a metaphorical interpretation of the entity CAN, will the immediate association between CAN and CONTENTS be sustained in the story.

Further more, the foregrounded meaning of CAN is highlighted in the first sentence "[w]e should all have a can like this" in Paragraph 6 (highlighted in Figure 5). In this sentence, the word "we" addresses generally to the reader, and the expression "a can" suggests the meaning of the CAN entity that is globally salient to the reader. The pronoun "this" foregrounds the information of the CAN and its CONTENTS which is locally salient in the story. By suggesting to the reader a possession of such a unique CAN, the textual information in Paragraph 6 foregrounds a contrast between the literal can and the metaphorical can, and invites the reader to perceive the "can" in a figurative sense.

The metaphorical meaning of CAN is not encoded in the reader's mental lexicon, but it is made available when the text proceeds from Paragraphs 3 to 6, and becomes the most salient information in Paragraph 6. However, although this newly developed meaning of "can" is salient in the local textual level, it will not become personally salient to the reader if it is rejected by the reader. This issue will be discussed in the following section.

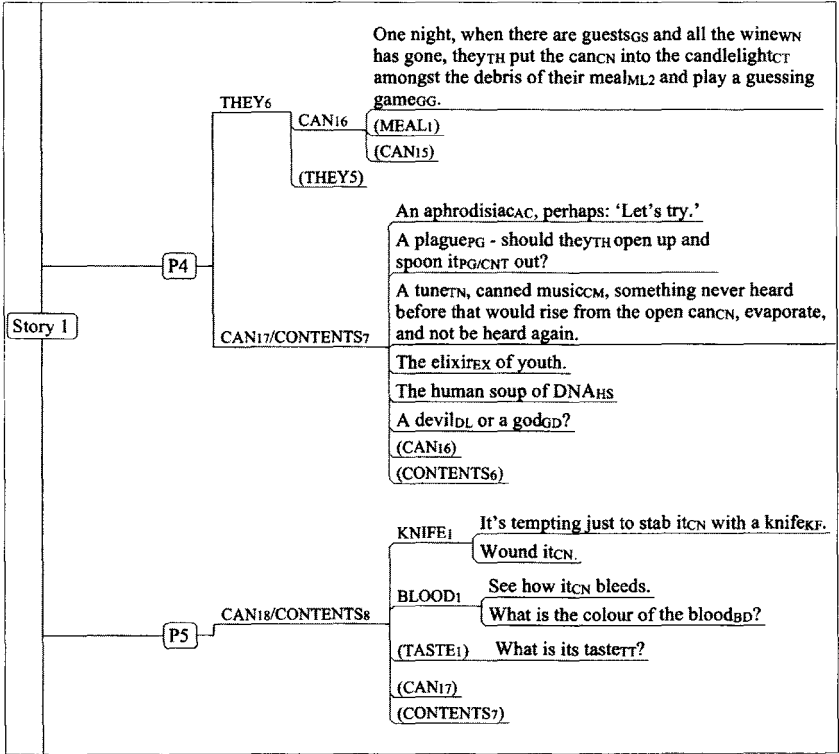


Figure 4. Textual information of *The Can Story*: Paragraphs 4–5

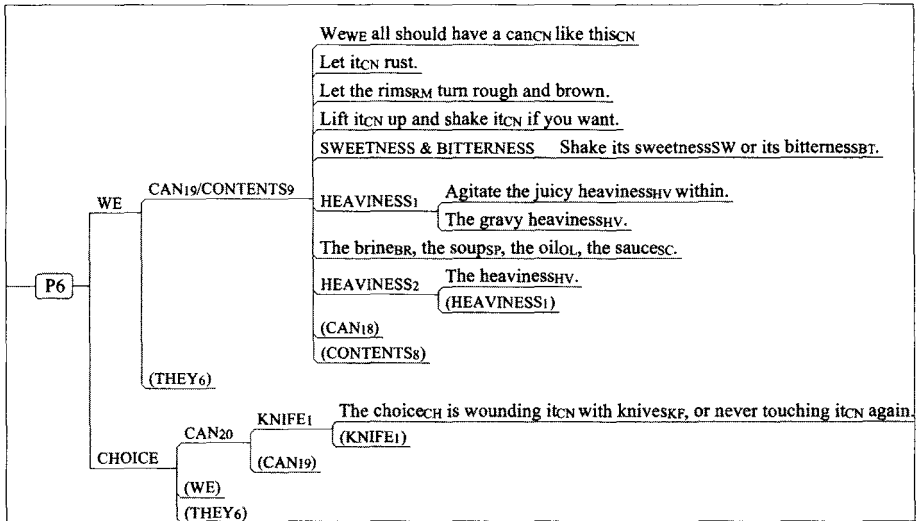


Figure 5. Textual information of *The Can Story*: Paragraph 6

4.4 The readers' preferred interpretations of *The Can Story*

In Huang (2008, Chapter 7), a small number of reader interviews are conducted in order to provide some insight to the readers' interpretive preferences. A total amount of eight readers participated in the interview. Each of the reader is instructed to read *The Can Story*, and is asked the same set of questions in regard to their interpretation of the story and a number of concepts relevant to the story.

In the initial reading stage, all eight readers interpret the entity CAN in its literal, food-related sense, which accords to the discussion of the salient encoded meaning of a mental entity at a global level. After the reading process, the metaphorical meanings of the can (such as "fear", "family secrets", and "the unknown") are available to all but one readers, and most readers develop the metaphorical reading from Paragraphs 3 and 4, which accords to the analysis of syntactic salience and foregrounded information which become apparent in the same location within the text. The results of the interview suggest the same information that is salient in the global level and in the local textual level is also accessible to the readers, and is likely to become personally salient to the readers as well. However, the readers can decide to sustain or reject this available information depending on their personal preferences. A reader who prefers facts to fiction is able to perceive a metaphorical meaning depicted in a literary narrative, but she may choose the literal interpretation to be the final interpretation if the text can be interpreted both literally and figuratively. In the interview of *The Can Story*, the literal meaning of the CAN sustains during the readers' reading processes, and co-exists with its metaphorical meaning in the readers' final interpretations of the story. This is probably due to two reasons, the first of which being that a literal interpretation of the story is possible; the second of which being the novel metaphorical meaning is only salient in a local textual level, and cannot completely override the literal meaning that is globally salient and deeply entrenched in the readers' mental lexicon. However, the interview result does show a reasonable amount of overlap between the textually salient meanings and the personally salient meanings in the process of *The Can Story* (see Huang 2008: 269–270). Among these meanings, some have gone through meaning consolidation and some have not, but each of them is highly accessible at a given time point of discourse processing.

5. Conclusion

This paper aims to investigate the development of metaphorical meaning in a discourse context. One way to solve this metaphorical riddle is to explain how certain information becomes more accessible than the other in the cognitive process of a metaphor. In order to tackle this issue, a notion of salience is proposed in this paper to examine dynamic accessibility of a meaning at a given time point of discourse processing.

This notion of salience is related to the one discussed in the Graded Salience Hypothesis and its two supplementary hypotheses (Giora 2003). Both notions of salience refer to the encoded meaning of a mental entity, but the newly proposed salience also includes four different types of information: the meaning with syntactic prominence; the meaning of an

entity associated with a salient entity; the foregrounded meaning in a discourse, and the reader's preferred meaning. These four types of meanings are highly accessible in the local and personal levels, and have the potential to become globally salient through a gradual development.

This proposed notion of salience is the basis of the salience-based model for metaphorical interpretation, which is an analytical framework that investigates the development of a metaphorical meaning in a discourse context. This paper has provided an analysis of a short story in order to demonstrate how salience dynamics can be monitored and examined from the aspects of the encoded meaning, the syntactic prominence, salient associates, and foregrounded information in a given discourse. The information that is globally and locally salient according to the analysis is further compared with the readers' interpretations in order to observe the reception and development of the metaphorical meaning from the readers' personal perspectives.

The notion of salience proposed in this paper, alongside with the salience-based framework developed in Huang (2008) and outlined in the previous sections aim to explain the construction of metaphorical meaning in a discourse context. Further study and research will be required to improve the current framework, and more empirical data are needed to support further research on salience dynamics in metaphorical comprehension.

Appendix

Text 1 extracted from Crace (2002), with added marks of paragraph information:

- (P1) Someone has taken off – and lost – the label on the can. There are two glassy lines of glue with just a trace of stripped paper where the label was attached. The can's batch number – RG2JD 19547 – is embossed on one of the ends. Top or bottom end? No one can tell what's up or down. The metal isn't very old.
- (P2) They do not like to throw it out. It might be salmon – not cheap. Or tuna steaks. Or rings of syruped pineapple. Too good to waste. Guava halves. Lychees. Leek soup. Skinned, Italian plum tomatoes. Of course, they ought to open up the can and have a look, and eat the contents there and then. Or plan a meal around it. It must be something that they like, or used to like. It's in their larder. It had a label once. They chose it in the shop.
- (P3) They shake the can up against their ears. They sniff at it. They compare it with the other cans inside the larder to find a match in size and shape. But still they cannot tell if it is beans or fruit or fish. They are like children with unopened birthday gifts. Will they be disappointed when they open up the can? Will it be what they want? Sometimes their humour is macabre: the contents are beyond description – baby flesh, sliced fingers, dog waste, worms, the venom of a hundred mambas – and that is why there is no label.
- (P4) One night, when there are guests and all the wine has gone, they put the can into the candlelight amongst the debris of their meal and play a guessing game. An aphrodisiac, perhaps: 'Let's try.' A plague – should they open up and spoon it out? A tune, canned music, something never heard before that would rise from the open can, evaporate, and not be heard again. The elixir of youth. The human soup of DNA. A devil or a god?

- (P5) It's tempting just to stab it with a knife. Wound it. See how it bleeds. What is the colour of the blood? What is its taste?
- (P6) We all should have a can like this. Let it rust. Let the rims turn rough and brown. Lift it up and shake it if you want. Shake its sweetness or its bitterness. Agitate the juicy heaviness within. The gravy heaviness. The brine, the soup, the oil, the sauce. The heaviness. The choice is wounding it with knives, or never touching it again.

References

- Anaki, D., Faust, M. & Kravetz, S. 1998. Cerebral hemispheric asymmetries in processing lexical metaphors. *Neuropsychologia* 36: 691–700.
- Carston, R. 2002. *Thoughts and Utterances: The Pragmatics of Explicit Communications*. Oxford: Blackwell.
- Cook, G. 1994. *Discourse and Literature: The Interplay of Form and Mind*. Oxford: Oxford University Press.
- Crace, J. 2002. *The Devil's Larder*. London: Penguin.
- Evans, V. Forthcoming. *How Words Mean*. Oxford: Oxford University Press.
- Fauconnier, G. & Turner, M. 2002. *The Way We Think*. New York: Basic Books.
- Gerrig, J. R & Rapp, N. D. 2004. Psychological processes underlying literary impact. *Poetics Today* 25(2): 265–281.
- Giora, R. 1997. Understanding figurative and literal language: the Graded Salience Hypothesis. *Cognitive Linguistics* 7: 183–206.
- Giora, R. 1999. On the priority of salient meanings: studies of literal and figurative language. *Journal of Pragmatics* 31: 919–929.
- Giora, R. 2002. Literal vs. figurative language: different or equal? *Journal of Pragmatics* 34: 487–506.
- Giora, R. 2003. *On our Mind*. Oxford: Oxford University Press.
- Giora, R. & Balaban, N. 2001. Lexical access in text production: on the role of salience in metaphor resonance. In *Text Representation*, T. Sanders, J. Schilperoord & W. Spooren (eds), 111–124. Amsterdam: John Benjamins.
- Giora, R. & Fein, O. 1999. On understanding familiar and less-familiar figurative language. *Journal of Pragmatics* 31: 1601–1618.
- Giora, R., Fein, O., Kronrod, A., Elnatan, I., Shuval, N. & Zur, A. 2004. Weapons of mass distraction: optimal innovation and pleasure ratings. *Metaphor and Symbol* 19(2): 115–141.
- Grosz, B., Joshi, A. K. & Weinstein, S. 1995. Centering: a framework for modeling the local coherence of discourse. *Computational Linguistics* 21(2): 203–225.
- Hajičová, E., Havelka, J. & Sgann, P. 2003. Discourse semantics and the salience of referents. *Journal of Slavic Linguistics* 11(1): 127–140.
- Hajičová, E., Kruijff-Korbayová, I. & Kruijff, G. M. 1998a. Salience in dialogues. In *5th International Congress of the International Association of Dialogue Analysis*, S. Cmejrková, J. Hoffmannová & O. Mülnerová (eds), 381–393. Prague: Max Niemeyer Verlag.
- Hajičová, E., Partee, B. H. & Sgall, P. 1998b. *Topic-Focus Articulation, Tripartite Structures, and Semantic Content*. Dordrecht: Kluwer Academic Publishers.
- Hajičová, E. & Sgall, P. 2001. Topic-focus and salience. *Meeting of the Association of Computational Linguistics and the 10th Conference of the European Chapter*, 268–273. CNRS: Morgan Kaufmann.
- Huang, M. Z. 2008. *Metaphor, Salience and Literary Discourse*. PhD dissertation, University of Nottingham, UK.
- Katz, A. & Ferretti, T. 2000. Reading proverbs in context: the role of explicit markers. *Discourse Processes* 36: 19–46.

- Katz, A. & Ferretti, T. 2001. Moment-by-moment reading of proverbs in literal and nonliteral contexts. *Metaphor and Symbol* 16: 193–221.
- Lakoff, G. & Johnson, M. 1980. *Metaphors We Live by*. Chicago: University of Chicago Press.
- Musolf, A. 2004. *Metaphor and Political Discourse: Analogical Reasoning in Debates about Europe*. Basingstoke: Palgrave MacMillan.
- Peleg, O., Giora, R. & Fein, O. 2001. Salience and context effects: two are better than one. *Metaphor and Symbol* 16 (3/4): 173–192.
- Pexman, P. M., Ferretti, T. & Katz, A. 2000. Discourse factors that influence irony detection during on-line reading. *Discourse Processes* 29: 201–222.
- Prentice, D. A., Gerrig, R. J. & Bailis, D. S. 1997. What readers bring to the processing of fictional texts. *Psychonomic Bulletin and Review* 4: 416–420.
- Rapp, D. & Gerrig, R. 2002. Readers' reality-driven and plot-driven analysis in narrative comprehension. *Memory and Cognition* 30(5): 779–788.
- Rapp, D., Gerrig, R. & Prentice, D. 2001. Readers' trait-based models of characters in narrative comprehension. *Journal of Memory and Language* 45: 737–750.
- Semino, E. 2004. Book review: *On Our Mind* (R. Giora, 2003). *Journal of Pragmatics* 36: 2185–2189.
- Sperber, D. & Wilson, D. 1995. *Relevance: Communication and Cognition*. Oxford: Blackwell.
- Turner, N. E. & Katz, A. 1997. Evidence for the availability of conventional and of literal meaning during the comprehension of proverbs. *Pragmatics and Cognition* 5: 203–37.
- van Peer, W. & Hakemulder, J. 2006. Foregrounding. In *The Pergamon Encyclopaedia of Language and Linguistics*, K. Brown (ed.), 546–551. Oxford: Elsevier.
- Walker, M., Joshi, A. K. & Prince, E. F. 1998. *Centering Theory in Discourse*. Oxford: Oxford University Press.

When is a linguistic metaphor a conceptual metaphor?

Daniel Casasanto

1. Introduction

In short, the locus of metaphor is not in language at all, but in the way we conceptualize one mental domain in terms of another. (Lakoff 1993: 203)

The central claim of Conceptual Metaphor Theory is that people *conceptualize* many abstract domains metaphorically, in terms of domains of knowledge that are relatively concrete or well-understood (Lakoff 1993; Lakoff and Johnson 1980, 1999).¹ George Lakoff (1993: 244) writes that “metaphor is fundamentally conceptual, not linguistic, in nature.” Yet, the overwhelming majority of evidence for conceptual metaphor *is* linguistic in nature. The linguistic data that can be marshaled in support of metaphor theory are compelling and varied. They include analyses of the systematicity of source domain – target domain relations (e.g. H. Clark 1973; Lakoff and Johnson 1980, 1999), patterns of semantic change throughout history (Lafargue 1898/1906; Sweetser 1991), patterns of child language acquisition (Bowerman 1994; Johnson 1999), computational modeling of abstract word meanings (Narayanan 1997), and experimental data on language processing (e.g. Boroditsky 2000, 2001; Gibbs 1994; Glenberg and Kaschak, 2002). But are linguistic data enough?

There are both *in principle* and *in practice* reasons why we cannot infer the structure and content of non-linguistic mental representations based solely on linguistic and psycholinguistic data. In principle, if Conceptual Metaphor is a theory of mental representation (and not just of language), then it must be true that people structure their abstract concepts metaphorically even when they're not using language. Yet, this claim is impossible to test with methods that require people to process abstract concepts *in language*. It is plausible that the mental representations people form when they are using language are importantly different from the mental representations they form when they are

1. Throughout this chapter, Conceptual Metaphor Theory will be used to refer to Lakoff and Johnson's proposal, as well as related theories. This generalization obscures some theoretical differences among proposals by different researchers, and even differences between Lakoff and Johnson's CMT circa 1980 and circa 1999. However, the present discussion should be equally relevant for all theories that attempt to predict the structure of abstract concepts based on patterns in metaphorical language.

perceiving, remembering, and acting on the world without using language (E. Clark 2003; Slobin 1987). Linguistic tests alone cannot evaluate this possibility. In practice, while some non-linguistic experimental results have validated Conceptual Metaphor Theory, others have challenged it. This chapter will briefly review experiments testing our understanding of the abstract domain of *time*, and then present three experiments exploring the metaphorical basis of *similarity*. In keeping with the ‘new directions’ theme of this volume, this chapter will illustrate how tools developed by psychologists can be used to test cognitive linguistic theories, and how experimental results can suggest novel conceptualizations of long-studied domains.

2. Time is our fruit fly

Time has become for the metaphor theorist what the fruit fly is for the geneticist: the model system of choice for linguistic and psychological tests of relationships between metaphorical source and target domains. Linguistic analyses evince intricate systems of conceptual projections from the source domains of *space* and *motion* to the domain of time (e.g. Alverson 1994; H. Clark 1973; Evans 2004; Grady 1997; Lakoff and Johnson 1980, 1999), some of which have been validated in psycholinguistic experiments (e.g. Boroditsky 2000, 2001; Torralbo, Santiago, and Lupiáñez 2006), gesture experiments (Casasanto 2008a; Núñez and Sweetser 2006), and low-level psychophysical tests (Casasanto 2008b; Casasanto and Boroditsky 2008). The latter experiments were designed expressly to address the *in principle* limitation of language-based studies, described above.

2.1 Spatializing time in language and thought

Our approach was to test whether the same asymmetric relationship between space and time found in linguistic metaphors is also found in people’s non-linguistic mental representations of time. People tend to talk about time in terms of space (e.g. a *long* vacation, a *short* engagement) more than they talk about space in terms of time (Lakoff and Johnson 1980). Do people also *think* about time in terms of space – more than the other way around – even when they’re not using language? To find out, Lera Boroditsky and I conducted a series of experiments in which people watched simple, nonverbal stimuli (e.g. a line ‘growing’ across the computer screen), and clicked the mouse to reproduce either the duration of the stimulus (i.e. how much time the line remained on the screen) or its spatial displacement (i.e. the distance of the line from end to end). Results showed the predicted space–time asymmetry. Participants could ignore a line’s duration when estimating its spatial distance, but they could not ignore distance when estimating duration. Lines that traveled a shorter distance were judged to take a shorter time, and lines that traveled a longer distance were judged to take a longer time – even though, in reality, all lines had the same average duration, regardless of the distance they traveled. Even when participants were warned which dimension of the stimulus they should pay attention to, they couldn’t help incorporating irrelevant spatial information into their temporal judg-

ments (but not vice versa). These experiments showed that the asymmetric relationship between space and time found in linguistic metaphors is also found in our more basic non-linguistic representations of distance and duration (Casasanto 2008b; Casasanto and Boroditsky 2008).

Subsequent experiments showed that relationships between non-linguistic representations of time and space are highly specific, and can be predicted based on particulars of a speaker's first language. Whereas English tends to use metaphors that liken time to spatial distance (e.g. 'a long time', like 'a long road'), other languages like Greek favor metaphors that liken time to an *amount of a substance* accumulating in three-dimensional space (e.g. *POLI ORA*, tr. 'much time', like 'much water').

English and Greek speakers performed a pair of psychophysical tasks to test how deeply linguistic metaphors might influence non-linguistic thought. The first task required them to estimate the duration of a growing line while ignoring its spatial length, as above (i.e. the *distance interference* task). The second task required them to estimate the duration of a container gradually filling up with liquid while ignoring its fullness (i.e. the *amount interference* task). English and Greek speakers showed strikingly different patterns of results. English speakers' duration judgments were strongly affected by line length, but only weakly affected by container fullness. Greek speakers showed the opposite pattern, as we predicted based on the relative strengths of the TIME IS DISTANCE and TIME IS AMOUNT metaphors in English and Greek. Training experiments showed that teaching English speakers to use amount metaphors for time in the laboratory caused them to perform the filling container task indistinguishably from Greek speakers.

These experiments suggest that linguistic metaphors not only reflect the structure of speakers' non-linguistic duration representations, they can also shape those representations (Casasanto 2008b). More importantly for the current discussion, they validate both the psychological reality of Conceptual Metaphor Theory and the *specificity* of the predictions it can make. We don't just think about time in terms of space, we think about time using exactly the type of spatial representations (i.e. linear or three-dimensional) that our linguistic metaphors imply. (See Boroditsky 2001; H. Clark 1973; Núñez and Sweetser 2006; Torralbo, Santiago and Lupiáñez 2006; and Tversky, Kugelmass and Winter 1991 for further evidence of the specificity of spatial schemas for time.)

Yet, despite such evidence supporting Conceptual Metaphor Theory, other tests have yielded conflicting results, even in the domain of time. Evans (2004) presents a catalog of discrepancies between the facts of English metaphors and the predictions that emerge from a recent, well-reasoned incarnation of metaphor theory, Grady's (1997) theory of Primary Metaphor. Evans points out that based on common English expressions like *we're coming up on the deadline*, the most natural inference is that English speakers mentally represent time in terms of upward motion on a vertical spatial axis. Yet, several lines of evidence (including Evans's informal survey of native speakers' intuitions about such statements) indicate that English speakers mentally represent events as if they follow one another along a horizontal spatial axis (Boroditsky 2000, 2001; H. Clark 1973; Núñez and Sweetser 2006; Torralbo, Santiago and Lupiáñez 2006; Tversky, Kugelmass and Winter 1991). It may be possible to address this particular concern of Evans's, in part, by pointing to analogous spatial expressions like *we're coming up to the front of the*

queue or *pull the car up to the curb* in which ‘up’ implies horizontal motion: this idiomatic horizontal use of ‘up’ occurs in the spatial domain as well as the temporal. As such, *coming up on the deadline* may arguably import a horizontal spatial schema into the domain of time. Still, the point remains that interpreting this spatio-temporal expression at face value would generate misleading predictions about the nature of non-linguistic time representations.

Trouble with time metaphors deepens when we consider other experimental results. Co-speech gestures corresponding to temporal expressions support Conceptual Metaphor Theory in some ways, but challenge it in others. Núñez and Sweetser (2006) interviewed Aymara speakers about how time expressions are used in their language, and then analyzed the gestures speakers produced during these interviews. They found that Aymara speakers often gestured forward when talking about the past and backward when talking about the future, consistent with the unusual spatial metaphors in their language that suggest the past lies ahead of them and the future behind them. By contrast, I conducted a series of experiments in which English speakers produced spontaneous co-speech gestures when telling stories about past and future events, but these gestures were largely inconsistent with spatio-temporal metaphors in English – and every other known language (Casasanto 2008a). English space-time metaphors place the future in front of the speaker (e.g., *the best years are ahead of us*) and the past behind the speaker, (e.g., *our salad days are behind us*), implying that time flows along the sagittal (front/back) axis. However, when English speaking participants told stories about sequences of events they systematically gestured along the transverse (left/right) axis, placing the past to the left and the future to the right (see also Calbris 1990; Cienki 1998; Núñez and Sweetser 2006). This was true whether they used spatial language explicitly (e.g., ‘a century before’) or expressed the same ideas using non-spatial language (e.g., ‘a century earlier’). These results are broadly consistent with the claim from Conceptual Metaphor Theory that English speakers mentally represent time in terms of horizontal space. Yet, they are inexplicable on a strict version of this theory given that left-right spatio-temporal metaphors are entirely absent from English speech. Cultural conventions such as reading, writing, and calendars that represent time as flowing from left to right point toward a partial explanation of this behavior (Tversky, Kugelmass and Winter 1991), but there is no obvious way to predict – or even account for – the left-right spatialization of time based on patterns in metaphorical language.

These spontaneous gesture data (Casasanto 2008a) not only raise questions about the relationship between linguistic metaphors and conceptual metaphors, they also challenge many English speakers’ intuitions about the way they gesture, and suggest a dissociation between people’s conscious and unconscious spatializations of time. When I asked English speakers informally to show how they typically gesture to indicate pastward and futureward events, they often gestured on the sagittal axis – placing the future in front of them and the past behind them, consistent with front-back metaphors in English. This was true for naïve informants and metaphor theorists, alike. Yet, these deliberate, conscious gestures (or gesture demonstrations) differ strikingly from the spontaneous gestures that experimental participants produced when they were not talking explicitly about the concept of time. Speakers’ conscious reflections on how they use space to represent time appear to be colored by the verbal metaphors at their disposal, but their unconscious representa-

tions of time reveal other non-linguistic sociocultural influences, as well. Understanding how space structures our mental representations of time will require integrating linguistic analyses and non-linguistic experimental results.

2.2 Time and speed

Speed also serves as a source domain for time in language, and provides another testbed for exploring the relationship between linguistic and conceptual metaphors. QUICKNESS acts as a metaphorical vehicle in utterances expressing either BREVITY (1a) or CONTRACTED DURATION (1b).

- (1) a. We'll take a quick vacation.
(QUICKNESS = BREVITY)
b. Our vacation went by quickly.
(QUICKNESS = CONTRACTED DURATION)

In the first example, the speaker comments on the duration of the vacation, *per se*, whereas in the second example the speaker comments that the duration felt contracted relative to expectation (whether or not the vacation was, in fact, brief). In both of these cases, speed is inversely related to duration, consistent with the relationship between velocity and time in Newtonian kinematics:

$$\text{Time} = \frac{\text{Distance}}{\text{Velocity}}$$

In this formula, time and distance are positively correlated, as suggested by metaphorical expressions like *a long party* and *a short concert*. The 'growing line' experiments reviewed above demonstrate that this positive correlation between distance and time exists in people's non-linguistic mental representations as well. Just as time and velocity are negatively correlated in this kinematic formula, time and speed (the scalar analog of velocity) appear negatively correlated in linguistic metaphors such as (1a) and (1b). Are time and speed also negatively correlated in people's non-linguistic mental representations?

Piaget's inquiries into children's understanding of time provide a surprising answer. Distance metaphors for time are similar in French and English (e.g. *depuis longtemps* means 'for a long time'). Consistent with these metaphors, Piaget found that French-speaking children often based their judgments of duration on their experience of distance. For example, when asked to judge the relative duration of two trains traveling along parallel tracks at different speeds, children often reported (erroneously) that the train traveling the longer distance took the longer time. Quickness metaphors for time also function similarly in French and English (e.g. *des vacances rapides* means 'a quick vacation' or 'a brief vacation'). Contrary to these metaphors, however, Piaget found that children often reported the train traveling at the *faster* speed took the longer time (Piaget 1927/1969; see also Mori, Kitagawa and Tadang 1974). Children believed that both distance and speed were positively correlated with time. Piaget concluded that time, space, and speed remain conflated in children's mental representations of motion events until about age nine, but

that after this age they construct the logical relationships among these dimensions suggested by Newtonian kinematics (and by linguistic metaphors).

Experiments by Casasanto and Boroditsky suggest the conflation of time, space, and speed in children's minds may be more enduring than Piaget realized. Adult English speakers from the MIT community performed a version of the growing line task that allowed the influences of distance and speed on time estimates to be evaluated independently. Our results were remarkably consistent with Piaget's. As in our previous studies, we found a positive relationship between distance and time: participants judged lines that traveled a shorter distance to take a shorter time, and lines that traveled a longer distance to take a longer time (even though, on average, all lines took the same amount of time, regardless of their spatial length). Surprisingly, we also found a positive relationship between speed and time: participants judged lines that traveled *slower* to take *less time*, and lines that traveled *faster* to take *more time* (even though, on average, all lines took the same amount of time, regardless of their speed). The effect of speed on time estimation was just as strong as the effect of distance on time estimation. This positive relationship between speed and time remained significant even when the influence of distance was removed mathematically, by partial correlation. This outcome was unexpected in light of the highly specific patterns of cross-dimensional interference observed in the space-time experiments described earlier, which were predicted from metaphors in participants' first languages. Based on these, we can rule out the possibility that participants simply construe more of one dimension in a motion event as more of another, indiscriminately. If that were the case, we would not have found the asymmetric interference between time and space described in the first set of growing line experiments, or the cross-linguistic differences in space-time interference patterns in the growing line/filling container experiments comparing English and Greek speakers.

Why did Piaget's children and our MIT undergraduates reveal mental representations of motion events in which time, speed, and distance were all positively correlated? Piaget, who was an associate of Einstein's, suggested a link between psychological time and the relativity of physical time. Yet, it is hard to imagine our primitive intuitions of time, space, and speed being shaped by something so counterintuitive as Einsteinian relativity. An alternative explanation invokes the intuitive physics of *projectiles*. Newtonian kinematics makes a host of simplifying assumptions that are violated by our everyday interactions with the physical world. Although the equation above shows an inverse relationship between time and (average) velocity, consider the relationship between these dimensions when a projectile is thrown with either greater or lesser force. When we throw a ball hard, it travels a longer distance, at a greater velocity, and for a longer time than when we throw it softly: Distance \propto Time \propto Velocity. By throwing and observing projectiles, we may learn that there are, in fact, positive correlations of time, speed, and distance in our everyday experience. These correlations may have given rise (in either evolutionary or developmental time) to the primitive, non-Newtonian understanding of time and speed revealed by Casasanto and Boroditsky's low-level psychophysical experiments and by Piaget's studies. Eventually, perhaps through language use and explicit instruction as well as through physical experience, children learn that under special circumstances time and speed are inversely related (e.g. when distance is held constant, as in many of Piaget's experiments,

or in the everyday experience of commuting from home to the office quickly or slowly along a given route).

This proposal, that the physics of projectiles shapes our intuitions of time, space, and speed, is speculative and in need of further investigation. What is important for the present discussion is that psychological tests reveal we have at least two contrasting ways of understanding the relationship between time and speed – only one of which can be predicted based on speed-time metaphors in language. Linguistic metaphors enshrine the more sophisticated inverse relationship between time and speed given by Newtonian kinematics. If our theory of how time is mentally represented were based solely on patterns in metaphorical language, we would never discover the more primitive relationship that governs children's understanding of time and speed, and influences ours as well.

3. Similarity and proximity: When does close in space mean 'close' in mind?

The domain of *similarity* provides another potential testbed for hypotheses about conceptual structure that are derived from linguistic metaphors. How do people judge the similarity of words, objects, or ideas? Despite concerns about its usefulness as a construct (Goodman 1972), similarity remains the focus of much psychological research, perhaps because our sense of similarity seems intimately linked with our capacity to generalize, to form categories, and to individuate concepts (Medin, Goldstone and Gentner 1993). In English (and many other languages), when speakers talk about similarity they often use words and expressions that describe spatial relations. Things that are similar along nearly any dimension can be described as *close*, and things that are dissimilar as *far*. For example:

- (2) a. These two shades of blue aren't identical, but they're *close*.
- b. The opposing candidates' stances on the issue couldn't be *farther apart*.

Is it possible that the way people talk about similarity reveals something fundamental about the way they conceptualize it? Our notion of similarity is abstract, like our ideas of *justice*, *love*, or *time* inasmuch as it is (a) vaguely and variably defined, (b) highly context dependent, and (c) mentalistic: lacking a concrete referent in the physical world that can be perceived through the senses. The experiments reported here tested the hypothesis that our notion of similarity depends, in part, on mental representations of physical distance (Casasanto 2008c). In three experiments, participants rated the similarity of pairs of words or pictures, which were presented at varying distances on the computer screen (i.e. close, medium, or far apart). A simple prediction was made based on the distance metaphors for similarity that are used in metric psychological models of similarity and in everyday language: if people think about similarity the way they talk about it (i.e. similar things are *close*), then participants should judge stimuli to be more similar when they are presented close together on the screen than when they are presented far apart.

3.1 Experiment 1: Abstract nouns

Experiment 1 tested whether participants would rate pairs of abstract nouns to be more similar in meaning when they appeared closer together on the screen. Abstract nouns (e.g. *Grief*, *Justice*, *Hope*) were chosen as stimuli for this first test of the relationship between similarity and proximity because the predicted influence of space on similarity may be most evident for similarity judgments about abstract entities that cannot be perceived directly through the senses.

3.1.1 *Methods*

3.1.1.1 *Participants.* 27 native English speaking participants from the Stanford University community performed this experiment, in exchange for payment.

3.1.1.2 *Materials.* 72 abstract nouns (concreteness rating < 400) between 4 and 10 letters long were selected from the MRC Psycholinguistic Database. Nouns were randomly combined into 36 pairs (e.g. *Grief-Justice*, *Memory-Hope*, *Sympathy-Loyalty*). Words were presented on an iMac monitor (1024 x 768 pixels resolution, 72 dpi) in 14 point courier font.

3.1.1.3 *Procedure.* Participants viewed word pairs in randomized order, one word at a time, and rated their similarity in meaning on a scale of 1 (not at all similar) to 9 (very similar). Before the first word appeared, a pair of empty 'picture frames' (150 pixels wide, 50 pixels high) appeared on the vertical midline of the screen for 500 ms. The centers of the frames were separated horizontally by 150 pixels in the Close condition, 300 pixels in the Medium condition, and 450 pixels in the Far condition. Pairs of Close, Medium, and Far picture frames appeared in one of four positions on the far left, middle left, middle right, or far right of the screen. This variation in position was orthogonal to the variation in distance between words, and was intended to reduce demand characteristics of the task. After 500 ms, the first word in each pair appeared for 2000 ms in the leftmost picture frame, then disappeared. After a 500 ms inter-stimulus interval, the second word of the pair appeared in the rightmost picture frame for 2000 ms. The words of each pair were presented serially rather than simultaneously to rule out low-level explanations for any observed differences in similarity ratings across conditions due to differences in saccadic activity or sharing of visual attention. Participants saw each word pair once, and the assignment of word pairs to conditions was counterbalanced across subjects.

3.1.2 *Results and discussion*

Results of Experiment 1 showed that stimuli were judged to be more similar when they were presented closer together than when they were farther apart (Fig. 1). Z-scored similarity ratings were compared using one-way ANOVA. Ratings differed significantly across conditions, both by subjects ($F(2,52) = 3.45, p < .04$) and by items ($F(2,105) = 4.49, p < .02$). A one-tailed paired-samples t-test showed a difference between Close and Far trials when analyzed by subjects (difference = 0.28, $t(26) = 2.22, p < .02$). A one-tailed independent-samples t-test confirmed this difference between Close and Far trials when analyzed by items (difference = 0.24, $t(36) = 2.74, p < .004$).




Figure 1. Similarity ratings for pairs of abstract nouns varied significantly as a function of their spatial separation on the screen. Pairs were judged to be more similar when they were presented closer together on the screen, consistent with predictions based on Conceptual Metaphor Theory. Error bars indicate s.e.m.

The finding that stimuli were rated more similar when presented closer together is consistent with predictions based on Conceptual Metaphor Theory. One concern in interpreting these results was that some of the word pairs were judged to have very low similarity in all conditions, and that the influence of proximity may have been restricted to these pairs for which word meanings were difficult to compare. However, when data were mean-split, the same qualitative relation between similarity and proximity was found for high-similarity and low-similarity pairs, analyzed separately.

3.2 Experiment 2: Unfamiliar faces

Experiment 2 tested whether the results of Experiment 1 would generalize to a different type of stimulus for which similarity had to be computed along different dimensions. To judge the similarity of the abstract nouns pairs, participants had to retrieve word meanings from memory, and to reason about unseen properties of abstract entities. Because the appearance of words is arbitrarily related to their meaning, the visual stimuli themselves provided little information (if any) that was relevant to the similarity judgment. Would distance still influence similarity judgments as in Experiment 1 even if more of the relevant information were given perceptually, in the visual stimuli themselves? According to Conceptual Metaphor Theory, it should.

Although ‘concrete’ entities that can be perceived directly are *not* posited to be structured metaphorically (Lakoff and Johnson 1999), people use the SIMILARITY IS PROXIMITY metaphor to describe similarity between both abstract and concrete things, alike: just as two abstract words can be said to be *close in meaning*, two lines can be *close in*

length, two paint chips can be *close in color*, two shirts can be *close in size*, and two faces can be *close in appearance*. The relationship between similarity and proximity in linguistic metaphors generalizes broadly (so broadly, in fact, that it is difficult to imagine a case in which similarity cannot be described in terms of distance). The same metaphor can describe similarity along both conceptual and perceptual dimensions. Therefore, if people conceptualize similarity the way they talk about it, the same prediction about the relation between similarity and proximity should hold for both conceptual judgments about abstract entities and perceptual judgments about concrete entities.

For Experiment 2, participants judged the similarity of pairs of unfamiliar faces. Whereas participants in Experiment 1 were instructed to judge similarity of abstract words based on their *meanings*, participants in Experiment 2 were instructed to judge similarity of faces based on their *visual appearance*.

3.2.1 Methods

3.2.1.1 Participants. 33 native English speaking participants from the MIT community performed this experiment, in exchange for payment.

3.2.1.2 Materials and procedure. 60 pairs of unfamiliar faces were constructed from a database of University of Pennsylvania ID card photos. Half were male-male and half were female-female pairs. Faces pairs were presented exactly as word pairs were presented in Experiment 1, with the following exception: the height of the 'picture frames' was changed to accommodate the size of the photos (150 pixels wide by 200 pixels high).

3.2.2 Results and discussion

Results of Experiment 2 showed that stimuli were judged to be more similar when they were presented *farther apart* than when they were presented closer together (Fig. 2). Z-scored similarity ratings were compared using one-way ANOVA. Ratings differed significantly across conditions, both by subjects ($F(2,64) = 3.61, p < .04$) and by items ($F(2,177) = 3.29, p < .04$). A two-tailed paired-samples t-test showed a difference between Close and Far trials when analyzed by subjects (difference = 0.16, $t(32) = 2.90, p < .007$). A two-tailed independent-samples t-test confirmed this difference between Close and Far trials when analyzed by items (difference = 0.12, $t(118) = 2.45, p < .02$).

Whereas in Experiment 1 closer stimuli were judged to be more similar, in Experiment 2 closer stimuli were judged to be *less* similar. Thus, Experiment 2 results not only fail to show an influence of proximity on similarity in the direction that was predicted based on Conceptual Metaphor Theory (i.e. closer = more similar), they also show a highly significant effect of proximity on similarity judgments in the opposite direction.

3.3 Experiment 3: Object pictures

Why did proximity have opposite effects on similarity ratings for abstract nouns and unfamiliar faces? Experiments 1 and 2 differed both in the kind of stimulus participants judged (i.e. verbal vs. pictorial) and in the kind of judgments they made (i.e. 'conceptual' judgments based on meaning vs. 'perceptual' judgments based on visual appearance). Ex-




Figure 2. Similarity ratings for pairs of faces varied significantly as a function of their spatial separation on the screen. Pairs were judged to be less similar when they were presented closer together on the screen, contrary to predictions based on Conceptual Metaphor Theory. Error bars indicate s.e.m.

periment 3 evaluated whether the results of Experiments 1 and 2 differed because of the type of stimulus or the type of judgment.

For Experiment 3, different judgments were made on the same set of stimulus pictures, which depicted common objects. Half of the participants were instructed to judge their similarity in visual appearance (a perceptual judgment), and the other half to judge their similarity in function or use (a conceptual judgment). If the difference between the results of Experiments 1 and 2 was due to a difference in the type of experimental materials used, then results of both Experiments 3a and 3b should resemble those of Experiment 2, in which pictorial stimuli were used: closer stimuli should be judged to be less similar, regardless of the type of judgment participants made. By contrast, if the difference between results of the first two experiments was due to participants judging abstract, unseen properties of the stimuli in Experiment 1 but judging concrete, perceptible properties of the stimuli in Experiment 2, then results of Experiment 3a (conceptual judgment) should be similar to those of Experiment 1 (i.e. closer stimuli should be judged more similar), whereas results of Experiment 3b (perceptual judgment) should be similar to those of Experiment 2 (i.e. closer stimuli should be judged less similar).

3.3.1 Methods

3.3.1.1 Participants. 40 participants performed Experiment 3a and an additional 40 performed Experiment 3b, in exchange for payment. All were native English speakers from the MIT community.

3.3.1.2 Materials and procedure. 30 pairs of objects were constructed from the Snodgrass and van der Wart line drawings. Objects were paired only within semantic categories (e.g. tools, clothing, furniture) to facilitate meaningful comparisons. Object pairs were presented as in previous experiments, with the following exception: stimuli appeared at one of two distances on the screen (instead of three), to maximize the difference between the Close condition, in which the centers of pictures were separated by 150 pixels, and the Far condition in which the centers of pictures were separated by 600 pixels.

3.3.2 Results and discussion

Results showed that during conceptual judgments (Experiment 3a), closer stimuli were judged to be more similar (Fig. 3, left). By contrast, during perceptual judgments (Experiment 3b), closer stimuli were judged to be less similar (Fig. 3, right). Similarity ratings were z-scored, and a mixed ANOVA with Distance (Close, Far) as a within-subjects factor and Judgment Type (Perceptual, Conceptual) as a between-subjects factor showed a significant 2-way interaction by subjects, ($F(1,178) = 12.23, p < 0.001$) with no main effects. This significant interaction was confirmed in 2-way ANOVA by items, with Distance (Close, Far) and Judgment Type (Perceptual, Conceptual) as between-subjects factors ($F(2,116) = 12.12, p < 0.001$), with no main effects.

Planned pair-wise comparisons tested the difference between Close and Far trials in Experiments 3a and 3b, by subjects and by items. Two-tailed paired samples t-tests showed that Close trials were rated significantly more similar than Far trials during conceptual judgments (Experiment 3a: difference = .10, $t(39) = 2.59, p < .02$ uncorrected, $p = .03$ after




Figure 3. Results of Experiment 3a (left) and 3b (right). Similarity ratings for pairs of object pictures varied significantly as a function of their spatial separation on the screen. For the same set of stimuli, the relation between similarity and proximity was consistent with predictions based on Conceptual Metaphor Theory during Conceptual Judgments (Experiment 3a), but inconsistent during Perceptual Judgments (Experiment 3b). Error bars indicate s.e.m.

Bonferroni correction), whereas Close trials were rated significantly less similar than Far trials during perceptual judgments (Experiment 3b: difference = .09, $t(39) = 2.46$, $p < .02$ uncorrected, $p = .04$ after Bonferroni correction) when analyzed by subjects. Two-tailed independent-samples t -tests confirmed that this same pattern was found when data were analyzed by items: Close trials were rated significantly more similar than Far trials during conceptual judgments (Experiment 3a: difference = .10, $t(58) = 2.35$, $p < .03$ uncorrected, $p = .04$ after Bonferroni correction), whereas Close trials were rated significantly less similar than Far trials during perceptual judgments (Experiment 3b: difference = .10, $t(58) = 2.56$, $p < .02$ uncorrected, $p = .03$ after Bonferroni correction).

An additional meta-analysis was performed, comparing the effect of distance on similarity ratings for Close vs. Far trials across Experiments 1, 2, 3a, and 3b. The effect of proximity on similarity judgments for each experiment was defined as the difference between participants' mean similarity ratings in the Close and Far conditions [Effect of Proximity on Similarity = (mean of normalized similarity ratings in Close condition) - (mean of normalized similarity ratings in Far condition)], and was compared across all experiments using one-way ANOVA ($F(3,136) = 8.81$, $p < 0.0001$; Fig. 4). Two-tailed pair-wise independent-samples t -tests showed significant differences between the effects of proximity on similarity ratings for Abstract Nouns vs. Perceptual Object Judgments (difference = .37, $t(56) = 3.28$, $p < .002$ uncorrected, $p = .01$ after Bonferroni correction), Abstract Nouns vs. Faces (difference = .44, $t(58) = 3.41$, $p < .001$ uncorrected, $p = .006$ after Bonferroni correction), Perceptual Object Judgments vs. Conceptual Object Judgments (difference = .19, $t(78) = 3.57$, $p < .001$ uncorrected, $p = .006$ after Bonferroni correction), and for Conceptual Object Judgments vs. Faces (difference = .26, $t(71) = 3.98$, $p < .0001$ uncorrected, $p = .0006$ after Bonferroni correction). Importantly, no differences were found between




Figure 4. Comparison of the effect of proximity on similarity ratings across experiments. Error bars indicate s.e.m. Closer stimuli were rated more similar during conceptual judgments (Experiments 1 and 3a, left columns) but less similar during perceptual judgments (Experiments 3b and 2, right columns).

the effects of proximity on similarity ratings for Abstract Nouns vs. Conceptual Object Judgments (difference = .18, $t(65) = 1.62$, ns) or for Perceptual Object Judgments vs. Faces (difference = .07, $t(71) = 1.14$, ns).

In summary, this meta-analysis shows that all pair-wise comparisons *between* judgment types (conceptual vs. perceptual) yielded highly significant differences, whereas pair-wise comparisons *within* judgment types yielded no significant differences: results of the two experiments requiring conceptual judgments differed from the results of the two experiments requiring perceptual judgments. By contrast, the results of the two conceptual judgment experiments did not differ from one another, and the results of the two perceptual judgment experiments did not differ from one another.

Together, the results of Experiment 3 and of the meta-analysis suggest that the contrasting effects of proximity on similarity judgments found for Experiments 1 and 2 were not due to superficial differences between the verbal and pictorial stimuli. Rather, the effect of proximity on similarity depends on the kind of judgment participants make: conceptual judgments about abstract entities or unseen object properties vs. perceptual judgments about visible stimulus properties.

3.4 General discussion of Experiments 1–3

Experiments 1–3 tested whether similarity ratings for words and pictures vary as a function of how far apart stimuli appear on a computer screen. Results showed that physical proximity influenced similarity judgments significantly in all experiments, but the direction of influence varied according to the type of judgment participants made. Closer stimuli were rated more similar during ‘conceptual’ judgments of abstract entities or unseen object properties (Experiments 1 and 3a), whereas closer stimuli were rated less similar during ‘perceptual’ judgments of the visual appearance of faces and objects (Experiments 2 and 3b). Conceptual judgments followed the simplest prediction based on the SIMILARITY IS PROXIMITY metaphor (Lakoff and Johnson 1999): when stimuli appeared closer in physical space they were judged to be ‘closer’ in participants’ mental similarity space, as well. Perceptual judgments showed the opposite pattern, however, contrary to predictions based on linguistic metaphors for similarity.

Can these results be accommodated within a Conceptual Metaphor framework? The outcome of Experiments 1–3 is broadly consistent with the claim that abstract entities are mentally represented metaphorically, whereas concrete entities that can be perceived directly are represented non-metaphorically, on their own terms (Lakoff and Johnson 1980, 1999). Still, Conceptual Metaphor Theory is hard-pressed to account for the difference between the effects of space on perceptual vs. conceptual judgments, given that the same spatial metaphors for similarity can be used to describe both low-level perceptual properties and high-level conceptual properties: similarities in appearance, function, or meaning can all be described using words like *close* and *far*. Thus, linguistic metaphors suggest that the same conceptual metaphor underlies our notions of both perceptual and conceptual similarity (see examples (2a) and (2b), above). Although Experiments 1 and 3a supported the metaphor-based prediction that stimuli presented closer in space would be judged to be more similar, Experiments 2 and 3b showed the opposite pattern of results. Overall

these studies pose a challenge to Conceptual Metaphor Theory, and suggest that we cannot necessarily infer relationships between similarity and proximity in people's non-linguistic mental representations from patterns in metaphorical language.

Previous studies have also reported positive associations between proximity and *conceptual* similarity for both abstract and relatively concrete entities. Sweetser (1998) observed that speakers sometimes bring their hands closer together in space to indicate the similarity of abstract ideas via spontaneous co-speech gestures. Goldstone (1994) asked participants to arrange various tokens of the letter "A" on the computer screen such that more similar tokens were positioned closer in space. Although in principle similarity between tokens of the letter "A" could depend on perceptual properties of the stimuli, Goldstone noted that when participants were asked to indicate similarity via spatial proximity they focused on "abstract commonalities" between tokens (1994: 385). Whereas participants' non-spatial same/different judgments of the "A" stimuli were driven by perceptual similarity, instructing participants to arrange stimuli according to the rule that 'closer = more similar' led them to "tap into a level of similarity that is relatively cognitive rather than perceptual" (*ibid.*). This complex relationship between spatial proximity, conceptual similarity, and perceptual similarity appears to have been unexpected in the Goldstone study, as it was in the present study.

Conceptual Metaphor Theory does not predict the pattern of data reported here, and it is possible that no current theory of similarity predicts it *a priori*. However, considering the computation of similarity to be a rational statistical inference based on regularities in our environment may help to situate the observed pattern of results in an ecological framework (Anderson 1991; Shepard 1987; Tenenbaum and Griffiths 2001). As Gestalt psychologists observed, the world appears to be pervasively clumpy (Wertheimer 1923/1938). Things that belong to the same category tend to be found close together, and also tend to be similar to one another compared with things that belong to different categories. Given that we are continually exposed to such organization, and that recognizing clumpiness may be useful for reasoning about our environment, it seems plausible that people implicitly learn and use a set of relations that could be called The Clumpiness Principle (building on Wertheimer's principles of proximity and similarity): Proximity a Similarity a Category Membership.

Tenenbaum and Griffiths (2001) proposed a Bayesian model according to which the similarity of two items is computed in terms of the probability that they are members of the same category (i.e. drawn from the same statistical distribution). In their model, the probability that items share category membership is proportional to the likelihood that they do given the information present in the stimuli, *per se*, and also proportional to the probability that they do given the observer's prior experience and stored knowledge. If we assume this generalization-based view of similarity, then in the present experiments participants' estimates of the probability that stimulus items belonged to the same category (and, therefore, of their similarity) depended in part on perceptible information given in the stimulus, and in part on their implicit knowledge of the Clumpiness Principle. In the case of *conceptual* similarity judgments, little relevant perceptual information was available in the stimulus items, so participants' heuristic use of the Clumpiness Principle was evident: greater proximity was used as an index of more probable shared

category membership and of greater similarity. In the case of the *perceptual* similarity judgments, however, participants' estimates of the probability that stimulus items belonged to the same category were likely to depend more strongly on the perceptible information given in the stimuli themselves, which overwhelmed any influence of the Clumpiness Principle.

On this proposal, when perceptible information was available in the stimuli (and was relevant to the task), participants used it. Participants may have judged closer stimuli to be *less* similar in Experiments 2 and 3b because proximity facilitates noticing small differences during perceptual judgments that might go unnoticed for stimuli presented farther apart.² By contrast, when perceptual information wasn't available in the stimuli (in Experiment 1) or wasn't relevant to the required judgment (in Experiment 3a), then participants' judgments reflected their heuristic use of the knowledge that proximity correlates with category membership and with similarity.

Thus, it may be possible to account for the contrasting effects of proximity on conceptual and perceptual similarity judgments if the computation of similarity is considered to be a process of rational inference that optimally combines perceptible information at hand with stored knowledge of experiential regularities (Anderson 1991; Shepard 1987; Tenenbaum and Griffiths 2001).

3.5 Summary of similarity and proximity experiments

Three experiments showed that changing the spatial separation between pairs of words or pictures on the computer screen changed the way people rated their similarity. Our notion of similarity appears to depend, in part, on our experience of spatial proximity, but not always as predicted by spatial metaphors in language. When participants made conceptual judgments about abstract entities or unseen object properties, stimuli presented closer together were judged to be more similar than stimuli presented farther apart, consistent with predictions based on linguistic metaphors. By contrast, when participants made perceptual judgments about visible stimulus properties, stimuli presented closer together were judged to be less similar than stimuli presented farther apart, contrary to predictions based on linguistic metaphors. These findings underscore the importance of testing Conceptual Metaphor Theory experimentally, and suggest that it is not possible to infer the relationship between similarity and proximity in people's non-linguistic mental representations based solely on patterns in metaphorical language.

2. Since all stimuli were presented serially this explanation requires that proximity still facilitates noticing small differences between stimuli even when members of a pair are never seen simultaneously. Although further research is needed to test this assumption, this seems plausible in light of research showing that the spatial location of visually presented information is automatically indexed in memory and accessed during retrieval, even when the spatial information is task-irrelevant (Richardson and Spivey 2000).

4. Conclusions

The studies reviewed here show both convergence and divergence between predictions based on linguistic metaphors and the results of behavioral experiments. Studies testing the conceptual metaphors TIME IS SPACE, TIME IS SPEED, and SIMILARITY IS PROXIMITY yielded some results that could not be predicted based on metaphors in English (or any known spoken language). Importantly, where these studies failed to support predictions based on linguistic metaphors they did not simply produce null effects. Rather, they provided clear evidence of relationships between source and target domains that were either orthogonal to the relationships encoded in language (in the case of the gesture experiments showing left–right spatialization of time), or directly contradictory to the relationships predicted by patterns in language (in the case of the time-speed experiments by Piaget and by Casasanto and Boroditsky, and also the experiments on perceptual similarity judgments reported here).

The first conclusion these studies support is that relationships between non-linguistic domains of knowledge cannot necessarily be inferred from metaphors in language. Linguistic metaphors reveal only a subset of the conceptual metaphors that appear to structure our mental representations of similarity and time. The second conclusion is that even when linguistic metaphors fail to predict the exact relationships revealed by behavioral tests, they nevertheless point to important links between the source and target domains. Space and time, speed and time, and proximity and similarity are not unrelated: rather, they appear to be related in more complex ways than linguistic analyses alone can discover. As such, linguistic metaphors should be treated as a source of *hypotheses* about the structure of abstract concepts. Evaluating these hypotheses – determining when a linguistic metaphor reflects an underlying conceptual metaphor – requires both linguistic and extra-linguistic methods, and calls for cooperation across disciplines of the cognitive sciences.

Acknowledgements

Thanks to Laura Staum Casasanto, Herb Clark, and Vyvyan Evans for helpful comments. This research was supported in part by an NSF Graduate Research Fellowship, by NRSA post-doctoral fellowship #F32MH072502, and by a grant from the Spanish Ministry of Education and Science #SEJ2006-04732/PSIC, DGI.

References

- Alverson, H. 1994. *Semantics and Experience: Universal Metaphors of Time in English, Mandarin, Hindi, and Sesotho*. Baltimore: Johns Hopkins University Press.
- Anderson, J. 1991. The adaptive nature of human categorization. *Psychological Review* 98 (3): 409-429.
- Boroditsky, L. 2000. Metaphoric structuring: Understanding time through spatial metaphors. *Cognition* 75 (1): 1–28.

- Boroditsky, L. 2001. Does language shape thought? Mandarin and English speakers' conceptions of time. *Cognitive Psychology* 43 (1): 1–22.
- Bowerman, M. 1994. From universal to language-specific in early grammatical development. *Philosophical Transactions: Biological Sciences* 346 (1315): 37–45.
- Calbris, G. 1990. *The Semiotics of French Gestures*. Bloomington: Indiana University Press.
- Casasanto, D. 2008a. Conceptual affiliates of metaphorical gestures. Paper presented at *International Conference on Language, Communication, & Cognition*. Brighton, UK.
- Casasanto, D. 2008b. Who's afraid of the Big Bad Whorff? Cross-linguistic differences in temporal language and thought. *Language Learning* 58 (177): 63–79.
- Casasanto, D. 2008c. Similarity and proximity: When does close in space mean close in mind? *Memory & Cognition* 36 (6): 1047–1056.
- Casasanto, D. & Boroditsky, L. 2008. Time in the mind: Using space to think about time. *Cognition* 106: 579–593.
- Cienki, A. 1998. Metaphoric gestures and some of their relations to verbal metaphoric expressions. In *Discourse and Cognition*, J.-P. Koenig (ed.), 189–204. Stanford: CSLI Publications.
- Clark, E. 2003. Languages and representations. In *Language in Mind*, D. Gentner & S. Goldin-Meadow (eds), 17–23. Cambridge, MA: MIT Press.
- Clark, H. H. 1973. Space, time, semantics and the child. In *Cognitive Development and the Acquisition of Language*, T. E. Moore (ed.), 27–63. New York: Academic Press.
- Evans, V. 2004. *The Structure of Time: Language, Meaning and Temporal Cognition*. Amsterdam: John Benjamins.
- Gibbs, R. W., Jr. 1994. *The Poetics of Mind: Figurative Thought, Language, and Understanding*. Cambridge: Cambridge University Press.
- Glenberg, A. M. & Kaschak, M. P. 2002. Grounding language in action. *Psychonomic Bulletin and Review* 9: 558–565.
- Grady, J. 1997. *Foundations of Meaning: Primary Metaphors and Primary Scenes*. PhD Dissertation, University of California at Berkeley.
- Goldstone, R. 1994. An efficient method for obtaining similarity data. *Behavior Research Methods, Instruments, and Computers* 26 (4): 381–386.
- Goodman, N. 1972. Seven strictures on similarity. In *Problems and Projects*, N. Goodman (ed.), 437–447. New York: Bobbs-Merrill.
- Johnson, C. 1999. Metaphor vs. conflation in the acquisition of polysemy: The case of *see*. In *Cultural, Typological, and Psychological Perspectives in Cognitive Linguistics*, M. Hiraga, C. Sinha & S. Wilcox (eds), 155–169. Amsterdam: John Benjamins.
- Lafargue, P. 1898/1906. *The Origin of Abstract Ideas: Inquiries into the Origin of the Idea of Justice and the Idea of Goodness*. Chicago: Charles Kerr and Co.
- Lakoff, G. 1993. The contemporary theory of metaphor. In *Metaphor and Thought*, A. Ortony (ed.), 202–251. Cambridge: Cambridge University Press.
- Lakoff, G. & Johnson, M. 1980. The metaphorical structure of the human conceptual system. *Cognitive Science* 4: 195–208.
- Lakoff, G. & Johnson, M. 1999. *Philosophy in the Flesh: The Embodied Mind and its Challenge to Western Thought*. Chicago: University of Chicago Press.
- Medin, D., Goldstone, R. & Gentner, D. 1993. Respects for similarity. *Psychological Review* 100 (2): 254–278.
- Mori, I., Kitagawa, O. & Tadang, N. 1974. The effect of language on a child's forming of spatio-temporal concept: On comparing Japanese and Thai children. *Science Education* 58 (4): 523–529.
- Narayanan, S. 1997. *Embodiment in Language Understanding: Sensory-Motor Representations for Metaphoric Reasoning about Event Descriptions*. PhD Dissertation, University of California, Berkeley.

- Núñez, R. & Sweetser, E. 2006. Looking ahead to the past: Convergent evidence from Aymara language and gesture in the crosslinguistic comparison of spatial construals of time. *Cognitive Science* 30 (3): 401–450.
- Piaget, J. 1927/1969. *The Child's Conception of Time*. New York: Ballantine Books.
- Richardson, D. & Spivey, M. 2000. Representation, space, and Hollywood Squares: Looking at things that aren't there anymore. *Cognition* 76: 269–295.
- Shepard, R. 1987. Toward a universal law of generalization in psychological science. *Science* 237 (4820): 1317–1323.
- Slobin, D. 1987. Thinking for speaking. Paper presented at the Annual Meeting of the Berkeley Linguistics Society, Berkeley, CA.
- Sweetser, E. 1991. *From Etymology to Pragmatics: Metaphorical and Cultural Aspects of Semantic Structure*. Cambridge: Cambridge University Press.
- Sweetser, E. 1998. Regular metaphoricity in gesture: Bodily-based models of speech interaction. Paper presented at the 16e Congrès International des Linguistes.
- Tenenbaum, J. & Griffiths, T. 2001. Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences* 24: 629–640.
- Torralbo, A., Santiago, J. & Lupiáñez, J. 2006. Flexible conceptual projection of time onto spatial frames of reference. *Cognitive Science* 30 (4): 745–757.
- Tversky, B., Kugelmass, S. & Winter, A. 1991. Cross-cultural and developmental trends in graphic productions. *Cognitive Psychology* 23: 515–557.
- Wertheimer, M. 1923/1938. Laws of organization in perceptual forms. In *A Source Book of Gestalt psychology*, W. Ellis (ed.), 71–88. London: Routledge.

Generalized integration networks

Gilles Fauconnier

The expression “*blends*” is often used to refer to a type of data where, very visibly, two or more inputs are partially mapped onto each other and selectively projected to a new mental space in which novel structure can emerge (Fauconnier and Turner 1994, 1998, 2002). Famous examples of such blends are *The Buddhist Monk*, *Regatta*, *Nixon in France*, *Complex Numbers*, *The Image Club*. As it turns out, far from being exceptional, marginal, or genre-specific, such blends are all over the place, and especially visible in fields as different as scientific discovery, humor, advertising, or religious rituals.

What warranted a new category for this kind of data when we first studied it was that it didn't fit into any of the known mapping schemes, in particular the source–target scheme of metaphor theory as understood at the time, or analogy, or metonymy, or simple framing.

Methodologically, the abundance of previously unnoticed (and hence never analyzed) “blending” data suddenly offered a wealth of empirical resources to study with precision the cognitive operations¹ of mapping and integration that made such blends possible. As the principles of conceptual blending became better understood, conceptual blending itself became a legitimate tool of discovery and analysis. Take for example the evolution of Lakoff and Núñez's work on the cognitive basis of mathematics, which became the excellent book *Where Mathematics Comes From*, published in 2000. This research started out as an exploration of the metaphorical underpinnings of mathematics, and it ended up offering powerful analyses of mathematical conceptualization in terms of conceptual blends (Lakoff and Núñez 2000).² A key cognitive construction discovered by Lakoff and Núñez, the “basic metaphor of infinity” was later correctly reanalyzed by Núñez as a conceptual blending template (Núñez 2005).

1. A cognitive operation is not the same thing as a brain mechanism. Blending is called a cognitive operation here in the same sense as other high-level cognitive operations such as analogy, metaphor, framing, recursion, viewpoint shift, etc. Little is known yet about how the brain carries out such operations. Interesting computational models designed to capture the properties of such operations and to be maximally compatible with neural architectures are examined in Feldman (2006).

2. See especially the sections on the number line, granulars and hyper-reals, the disk-line segment, negative numbers, and the entire part VI of the book, which gives a superb description of many successive conceptual blends in mathematics that end up giving meaning to the mysterious formula $e^{i\pi} = -1$.

The systematic study of *integration* as a cognitive operation made many useful descriptive distinctions possible. So, within the data referred to as “blends”, there are different products depending on the types of inputs, the links between them, the choices for projection, etc. Corresponding types of blends are distinguished, or rather aligned on a graded continuum, going from simplex blends to mirror blends to single-scope and double-scope blends, all dividable into further subcategories.³ Blends can also be classified along other dimensions for various purposes.⁴

The description and classification of this new data is pretty much uncontroversial and widely viewed as innovative and useful. But a deeper project is to explore the role of integration and compression in meaning construction beyond the very visible blends that brought these cognitive operations to our attention.

In the present paper, I will point out some useful generalizations that emerge from the study of integration, along with some of the pervasive fallacies that stand in the way of making such generalizations. Through the analysis of attested data, I will discuss the notion of “generalized integration networks” and how they allow the construction of a multiplicity of surface products in human thought and action.

1. Three fallacies

Fallacy 1: Different surface products result from different cognitive operations

When the new surface products known as “blends” were first displayed, they were contrasted with familiar surface products and mapping operations that have been around forever: frames and framing, metaphors and metaphorical mappings, “logical” counterfactuals and counterfactual mappings, analogies and analogical mappings. Common sense and intellectual tradition associate a specific type of mapping with each specific type of surface product. Metaphorical mappings produce metaphors, analogical mappings produce analogies,⁵ and so on. In fact, the words “metaphor”, “analogy”, “metonymy”, are ambiguous: they can refer to the surface product or to the mapping that supposedly produced it. So it seems to make sense to see “blends” as one additional type of surface product, and to see “blending” as the specific mapping pattern associated with the newly discovered product.

This leap from product to process is fallacious. There is usually no isomorphism between the surface products as we see and classify them in everyday life and the underlying principles that produce them. Hence, once we have the evidence for integration (an operation) as provided by blends (data), it’s an open question whether that operation is restricted to “blends” or whether it might also be at work in more familiar data.

3. The taxonomy of blends is developed in Fauconnier and Turner (1998, 2002).

4. Some possibilities are outlined in Fauconnier and Turner (1994).

5. A pair of situations is not an analogy in any absolute or a priori sense. It becomes one if an analogical mapping is imposed on partial mental models for the situations.

Fallacy 2: If it's new, it's going to cost more

This is a triple fallacy, which applies equally when by “new” we mean “newly discovered” or when we mean “newly acquired” (through evolution or through learning).

The automatic assumption is that the recently noticed data (blends in the case at hand) must somehow be more exotic, less typical, than the familiar products already comfortably, if not always neatly, categorized. And along with that assumption also comes the corollary that more exotic phenomena demand extra cognitive effort, special machinery not normally used but available for out of the way (or even perhaps outlandish) human behavior in exceptional (marked) genres such as sarcasm, humor, spirituality, mathematics.

But this reasoning is also fallacious. New (i.e. previously unnoticed) data may indeed reveal cognitive operations; it does not follow that such operations are confined to the new data (Fallacy 1), or that they are necessarily exceptional, atypical, or costly. The data is new to the scientist, but it's always been around. In the case of blends, the data provides solid evidence for the operation of conceptual integration and the general principles that govern it.⁶ Humans happen to possess this capacity, including its double-scope manifestation, arguably indispensable for many singularities of human behavior (language, religion, science, art). There is no reason to believe that using this capacity is cognitively costly for humans. In fact, it turns out that humans use it all the time, for better or for worse, and that they clearly enjoy activities that depend on it, such as humor, deceit, rituals, or fiction.

A second corollary of Fallacy 2 is that existing theoretical machinery (e.g. standard framing or source–target metaphor theory) should be used whenever possible and that “new” machinery should only be resorted to when all else fails. In other words, the fallacious idea that less familiar data is cognitively more costly to produce is mirrored by the equally fallacious idea that “new” theoretical mechanisms (conceptual integration in this case) are theoretically more onerous than the more familiar theoretical mechanisms.

Moreover, Fallacy 2 is applied to learning: a child (it is assumed) learns the “simple” operations first and the “complex” ones later. Cognitively costly operations are acquired later, according to this implicit extension. Again, there is no a priori reason to believe this. Eighteen-month-old children (and perhaps younger ones) produce and understand elaborate blends.

Fallacy 3: A wide-ranging cognitive operation purports to explain “everything” Corollary: Such an operation explains “too much” and is unconstrained

It is generally agreed that “visible blends” reveal the mechanics of integration more readily than other phenomena. The obvious question is “does integration play a role in phenomena other than visible blends?” Considerable attention has been lavished on this issue in the last twelve years, with unexpected but (with hindsight) unsurprising convergent results: yes, integration operates in the construction of a multitude of surface products,

6. Constitutive principles, governing optimality principles, and compression laws are discussed in *The Way We Think* (Fauconnier and Turner 2002).

which may differ sharply from one another along many other dimensions. These surface products include for example analogical counterfactuals, mathematical notions like complex number, technological innovations like timepieces or computer interfaces, and also more familiar products like frames, metaphor, and grammatical constructions.⁷

Jerome Feldman, in his excellent work on the neural theory of language, writes:

More recently, Mark Turner and Gilles Fauconnier (2002) have made a bold attempt to explain much of mental life in terms of the cognitive linguistic notion of conceptual integration (or blending) we discussed in chapter 24. (Feldman 2006)

The proper way to understand this comment is that we (Turner and Fauconnier) have indeed shown that conceptual integration plays a necessary role in human mental life as evidenced by surface products of particular interest to humans.

But “necessary” is not “sufficient”. Billions of years of biological evolution precede the appearance of fully-fledged double-scope integration. Integration is only a minuscule component of the stunningly complex organization of the embodied mind. To be sure, it is responsible for striking singularities which distinguish humans from other species in ways that are of particular interest to humans themselves, and therefore to human researchers.

2. An example: The smoking ears network

When we avoid Fallacies 1, 2, 3, we can look at any surface product in a more general way. Instead of trying to fit it into a conventional descriptive category (such as metaphor, counterfactual, etc.), we can look in detail at the succession of mappings and integrations that operate in order to yield the complete surface product. Typically, what we find is a generalized integration network, which combines conventional integrations available in the language and culture with novel integrations and emergent structure made possible by the context in which the surface product is constructed.

To illustrate this, I will discuss an anecdotal example in some detail. The piece of data is taken from a column in *the San Francisco Chronicle* titled ‘Bar Patrons Fume Over Smoking Law’. The occasion for the newspaper column is the enforcement of the ban on smoking in the state of California.

“No Smoking” signs were tacked up in bars all over California yesterday, and hard-core smokers nursing a scotch or a beer were so angry that if they had been allowed to light up, the smoke would have been coming out of their ears.

2.1 The anger network

The excerpt from the Chronicle is immediately and effortlessly understood by readers, even though it triggers the construction of an elaborate “generalized” integration network.

7. See the extensive bibliography for such studies at <http://blending.stanford.edu>, and representative work in Coulson and Oakley (2000, 2005).

To build the network, a reader must have available the well known conventional metaphorical network of “anger as heat in the body container”. It has three major inputs, heat in a container, emotion, and body. I quote below from *The Way We Think* (Fauconnier and Turner 2002) in order to sum up the main features of this network. Kovecses and Lakoff’s analysis of the heat/anger metaphor is discussed in Lakoff (1987).

We have independently manipulable spaces for the emotion of anger and bodily states. We also have a conventional cultural notion of their relationship, based on correlation – people often do get flushed and shake when they are angry. We will call this notion the “Story of Emotion and Body”.

In addition to the metaphoric mapping between Heat and Emotions and the vital relation connection between Emotions and Body, there is a third partial mapping between Heat and Body. In this mapping, steam as vapor that comes from a container connects to perspiration as liquid that comes from a container; the heat of a physical object connects to body heat; and the shaking of the container connects to the body’s trembling.

The three partial mappings set the stage for a conventional multiple blend in which the counterparts in the inputs are fused, giving, for example, a single element that is heat, anger, and body heat, a single element that is exploding, reaching extreme anger, and beginning to shake. Once we have this blend, we can run it to develop further emergent structure, and recruit other information to the inputs to facilitate its development.

For example, we might say, “He was so mad I could see smoke coming out of his ears”. This derives from recruiting ears to the Body input and an orifice to the Heat input, and projecting them to the same element in the blend. We now have a new physiological reaction – smoke coming out of the ears – that is inconceivable in the original Body input. In the blend, it is fused with anger. Conventional expressions like “He exploded” can also prompt for new physiological reactions in the blend that are impossible for the Body input itself. In these cases, the notion of physiological correlates of emotion is coming from the “Story of Emotion and Body” inputs, but the specific content of the physiological reaction (smoke, explosion) is coming from the Heat input. This is a Multiple-Scope Network, with a conventional global generic space (Story of Emotions and Body) over two of the inputs and their vital relations, and with a systematic compression of those outer-space vital relations to uniqueness in the blend.

The blend remains linked to the inputs. A sentence like “He was so mad I could see smoke coming out of his ears” directly identifies structure in the blend, but inferences – smoke is a sign of great anger – are projected back to corresponding inferences in the Emotion input and the Body input: he was extremely angry and was showing physiological signs of it. (What these signs actually were in the actual human situation is irrelevant.)

Expressions can refer directly to the blend, as in “He exploded. I could see the smoke coming out of his ears.” This description, which would be inappropriate for any of the input spaces by itself, coherently picks out the integrated scene of the blend. Additionally, even when the vocabulary is appropriate for one of the input spaces, the blend can often use it in ways that would be ungrammatical for that input: for example, suppose the chef is angry and acts it out by boiling a pressure cooker until it explodes; although “anger” and “explode” apply to this scene, and although we could say the cooker “exploded with force”, we cannot say it “exploded with anger”. But in the blend, where the anger is pressure and heat and force, we can indeed say, “He exploded with anger”.

Vocabulary from all three inputs can be combined when referring to the blend, as in “She became red with anger and finally exploded”. Again, however, we could not say of a pan heated red by the angry chef that it was “red with anger”.

Running the blend can produce elaborate emergent structure, as in “God, was he ever mad. I could see the smoke coming out of his ears – I thought his hat would catch fire!”

There are no burning hats in the heat input or in the anger input. Burning hats are emergent in the blend, which has the frame of somebody on fire. They imply greater heat/anger, greater loss of control, and greater danger. (Fauconnier and Turner 2002: 300–301)

Table 1. Inputs to the anger network

Heat input	Emotion input	Body input
“physical events”	“emotions”	“physiology”
container	person	body
substance/liquid		blood
pressure	degree of anger	blood pressure
heat	anger	body heat
steam	sign of anger	perspiration, redness
explode	show extreme anger	acute shaking, loss of physiological control
boiling point	highest degree of emotion	
orifice (lid, spout, ...)		orifice (ears, navel, mouth)

From *The Way We Think* (Fauconnier and Turner 2002: 300).

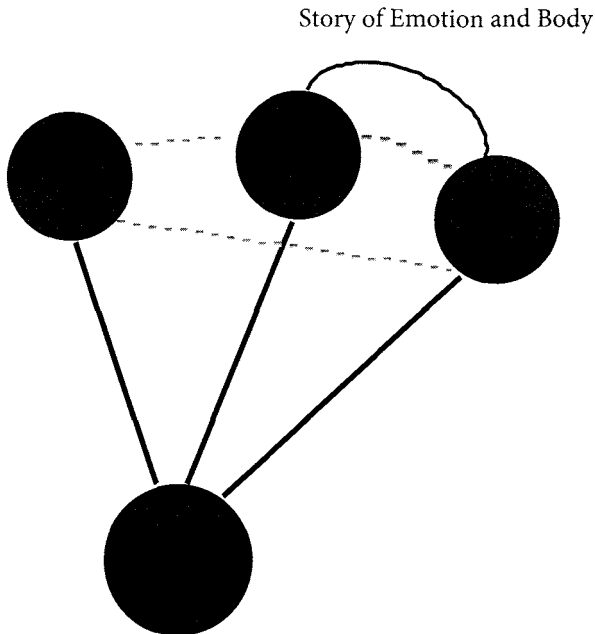


Figure 1. Anger network

2.2 Counterfactual “zoloft” networks

The newspaper column takes for granted the reader’s mastery of the culturally sanctioned anger network, and proceeds to build a counterfactual *if they had been allowed to light up* This is not a conventional network, but it conforms to a more general counterfactual blending template in which the blended space seems to incorporate incompatible pieces of information. Such networks are used fluently by subjects in many situations, and have been discussed by Coulson (2001) in the case of “abortion rhetoric”. I am calling them “zoloft” networks for present purposes, because of the following typical example of such networks provided by the so-called Zoloft defense case.⁸

In the Zoloft defense case, a teenager had murdered his grandparents. When the time came for him to be sentenced, his father pleaded for leniency and invoked the grandparents his son had murdered, saying: “if they were still alive, they would also plead for mercy for their grandson.”

Two inputs are blended: one in which the killing occurs, and the grandparents are dead, and one (counterfactual) in which nothing happens. In the latter, the grandparents are alive, they love their grandson as always, and he has done nothing wrong.

In the diagram below, *s* refers to the son/grandson, *g* to the grandparents.

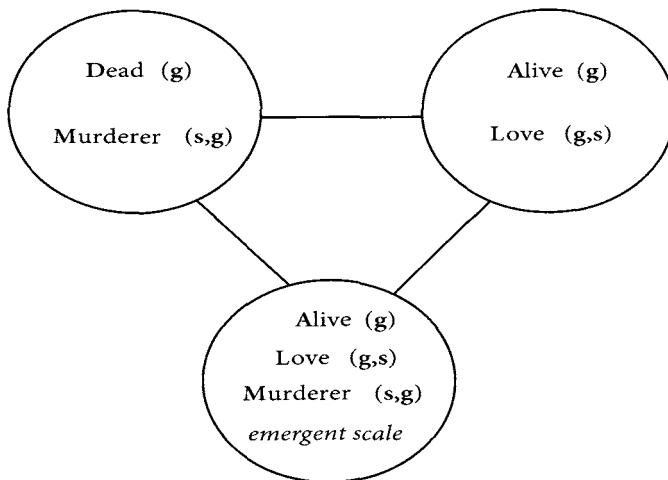


Figure 2. Zoloft network

8. The defense in this case was called the “Zoloft defense” because it argued that the grandson under the influence of the anti-depressant Zoloft, was not aware of the import of his actions. This defense failed in large part because the defendant, in addition to stabbing his grandparents in their bed, had later set fire to their house to destroy evidence of his crime.

The construction of the network is fairly straightforward. Some elements are brought in from the “reality” input: the grandson has been found guilty of murder and is being sentenced. Other elements are brought in from the counterfactual input: his grandparents love him and wish him the best. Emergent structure in the blended space is created in interesting ways. First, since the grandparents love their grandson and think highly of him, it follows logically that they plead for leniency. Secondly, a background cultural frame is covertly activated:⁹ pleas in favor of a convicted criminal carry different weight depending on who makes them. At the very top of the scale of credible advocates for the criminal (or indeed more generally for a culprit of any sort) is the victim of the crime. A victim can offer forgiveness, and/or plead effectively in favor of the offender. The underlying folk model includes a scale on which it is normally the victim who demands the harshest punishment, so that the actual punishment should not be more severe than what the victim demands. In the blended space of the zoloft network, the grandparents’ support for their grandson in ordinary circumstances becomes the dead victims’ support for the convicted murderer. The blend allows the scale to emerge with the dead grandparents at the very top.

The blended space seems of course “illogical”: if the grandparents were alive, there would be no murder, and therefore no sentencing. But as in many other blends that we use routinely, this apparent incongruity is not an obstacle to the emergence of the desired scale and the argument in favor of leniency that it provides.

2.3 Smokers’ zoloft network

The same zoloft blending template is used in the Fuming Smokers column. From one input NS (reflecting the real situation at hand: no smoking), we project the smoking ban (analogous to the murder in Zoloft), and the anger it triggers, and from the other (counterfactual) input YS, in which smokers smoke, we project the fact that they smoke, in order to obtain an emergent form of their anger, at the top of a scale of angry reactions. In the blended space Z of this zoloft network, the ban on smoking is in place causing smokers to be enraged, and they are allowed to smoke.

2.4 Angry smokers’ network

However, all this is not enough to account for the observed data. In the conventional metaphorical blend of anger, we do have the property that the greater the heat, the more steam, smoke or fumes will come out of orifices in the container, and in the blended space this yields the emergent property that great anger/heat in the body/container will cause steam/fumes to come out of orifices in the body/container: *He was fuming/steaming (with anger). He was so mad, you could see the smoke coming out of his ears.*

But this metaphorical smoke is not tied to any actual smoke in the body. What happens in the Fuming Smokers column is that a novel mapping is created opportunistically:

9. This is a standard process in the emergence of novel structure in the blended space, as shown for example in the analysis of stock examples like *The debate with Kant*, or *Regatta*.

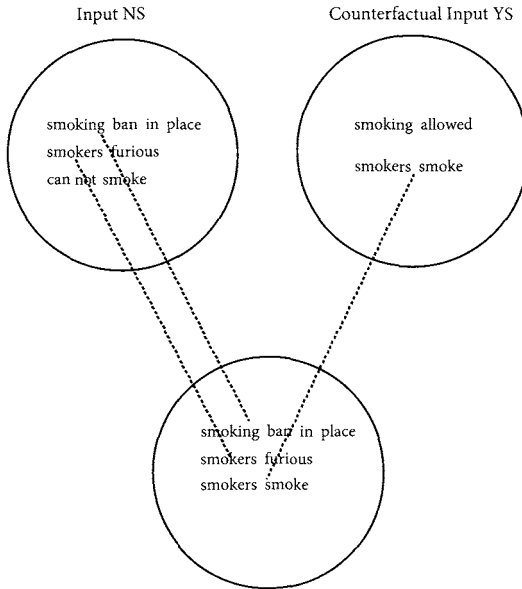


Figure 3. Smokers' zoloff network

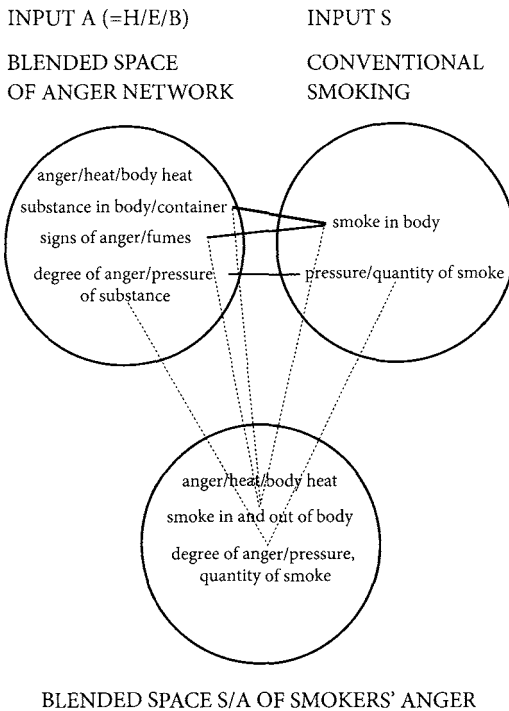


Figure 4. Angry smokers' network

the fumes from the anger network can be mapped onto the smokers' smoke. This mapping triggers a novel integration, the SMOKER'S ANGER network in which the smoke in the smokers' bodies is fused with the metaphorical fumes from the smokers' anger. One input is the blended space of the conventional ANGER network, where metaphorical heat is fused with body heat and with the corresponding emotion, anger; the other input is the frame of SMOKING, with a person inhaling smoke into their body. Furthermore in the blended space of this new integration, smoke is now the substance contained in the smoker's body. In the input of SMOKING, the smoke inside the smoker's body is not the result of some other substance heating up within the body; it is produced by an external heated object (burning tobacco) and *then* inhaled and exhaled by the smoker. In the blended space of the SMOKER'S ANGER network, the substance under pressure in the body and the smoke emitted as a result are fused. This is a *cause-effect* compression.

In the blended space constructed through this elaborate process, the smokers display their anger in a supremely visible way: we see the smoke from their burning tobacco coming out of orifices in their body.

But the situation described in the newspaper is one where a ban on smoking has just been put in place. This makes smokers angry but prevents their anger from being seen since they have no smoke to work with. The zoloft network Z (described in 2) comes to the rescue: it allows the construction of a further counterfactual blend Z/SA in which the ban is in place, and yet the smokers can express their anger as specified by the SMOKERS' ANGER network (i.e. by expelling tobacco smoke through their ears).


This is achieved by blending Z (the zoloft blended space in which the ban is in place, the smokers are angry, and they are allowed to smoke) with SA, the ANGRY SMOKERS blended space. In Z/SA, the ban is in place, and the smokers' rage can manifest itself at the top of the anger scale: smoke coming out of the ears.

2.5 Full network

We end up with the full network diagrammed below. The H/E/B part of this network is the conventional culturally sanctioned anger network, with a metaphorical component (anger as heat) and a metonymic component (anger as its stereotypical bodily manifestations). The integration is successful because H (heat, pressure, and substance in a container) also maps naturally onto B (body temperature, blood and blood pressure) in a way that is neither metaphorical nor metonymic.

H/E/B (conventional anger) is blended with S (smoking) to yield SA (smokers' anger). This part of the network is not conventional: it is creative in context, drawing opportunistically on a contextually available correspondence between the metaphorical smoke in H/E/B and the real smoke produced by the activity of smoking. Moreover, its only purpose is to feed into the next part of the generalized integration network, the zoloft network, in order to convey the thwarted smokers' fury.

That sub-network is built by using the available zoloft template (integrating an actual input with its counterfactual alternative). The input space NS, in which the ban is in place, causing smokers to be furious and preventing them from smoking, is blended with the counterfactual input space YS, in which they are allowed to smoke. In the blended space



H: heat	SA: smokers' anger
E: emotions	NS: no smoking space
B: body	YS: smoking allowed (counterfactual)
H/E/B: anger blend	Z: smokers's zoloft space (smoking ban+smoke)
S: conventional smoking	SA/Z: thwarted smokers fume

Figure 5. Full network

Z of the zoloft sub-network, the ban is in place, smokers are furious (projection from NS), and they smoke (projection from YS).

Z can now map onto SA, since both are cases of furious smokers who smoke. In the resulting blend, Z/SA, the ban is in place (projection from Z) and the smokers' fury is manifested by smoke coming out of their ears (projection from SA).

3. Discussion

The surface product exhibited in the newspaper excerpt is the result of several successive integrations. Is it a "metaphor"? Clearly, yes, but not a conventional one, and not simply a source to target mapping. Is it a "counterfactual"? Clearly, yes, since it builds on a counterfactual situation where smokers can smoke in spite of the ban, but it is not a case of building an alternative possible world. In possible worlds, smoke does not come out of the ears of smokers even when they are supremely angry. In possible worlds, smoking cannot be simultaneously banned and allowed. Is the newspaper statement contradictory or unintelligible? Clearly, no. On the contrary, it cleverly conveys the writer's point: that smokers are furious, frustrated, and unhappy. Is the piece of data a "blend"? Clearly, yes, and indeed a very visible one.

So, this surface product cannot be classified in a single category, or linked with a single mapping pattern. Understanding it requires the careful study of the multiple integrations that operate and of the established networks or network templates that are recruited effortlessly for its construction. It is not a “blend” as opposed to a “metaphor” or a “metaphor” as opposed to a “counterfactual”. It shares features with all of them.

Is attested data of this kind cognitively significant? Of course. Tens of thousands, perhaps millions, of readers of this newspaper column were able to perform the meaning construction with no conscious effort, no puzzlement, and presumably with pleasure, all of which attests to the universality of the processes involved. Is the cognitive construction a special one, to be distinguished from “ordinary” semantics? Certainly not, since as shown in the analysis it makes use of completely standard integration capacities, and recruits culturally entrenched networks and network templates.

Is the capacity for building generalized integration networks restricted to specific genres like humor? Solid evidence to the contrary has now been advanced in countless studies by scholars looking at conceptual blending. They have shown for example that elaborate generalized networks of the type described here are culturally and psychologically elaborated in the evolution of mathematical concepts, the technological development of instruments, of computer interfaces, of grammar for signed and spoken languages, of literary creativity in poetry, theatre, and literature, of courtroom practice, of religion and magic, or again of culturally significant and widely shared notions like the concept of “time”.¹⁰

This very rich body of work, like the more accessible “smokers’ fury” discussed here, dispels the fallacies mentioned at the outset of this paper. Using generalized integration, a capacity available to humans, is not more costly or exceptional than using other capacities, and simply interacts with them (fallacy 2). Integration is not a “theory of everything” (fallacy 3): it is a precise meaning construction operation that shows up in very diverse human products and exists in addition to, and not instead of, the infinitely greater body of biological and cognitive capacities available to us and often shared to some degree with other species. The surface distribution of products that we distinguish sharply in our everyday experience does not correlate with a corresponding distinction between the mapping schemes or cognitive capacities needed to produce them (fallacy 1): an operation like integration can be at work in all of them, and precise analysis of the generalized networks needs to be done for any observed data, just as chemical analysis needs to be done for any unknown chemical, without adding new elements or new principles to chemistry.¹¹

10. For the role of blending in the evolution of mathematical concepts, see Fauconnier and Turner (2002), Lakoff and Núñez (2000), Robert (1998). For the technological development of instruments, see Hutchins (2005), Williams (2005), Alac (2006). Blending and computer interfaces are discussed in Fauconnier (2001), Imaz and Benyon (in press). The role of blending in grammar for signed and spoken languages is demonstrated in Liddell (1998, 2003), Mandelblit (1997). Literary creativity in poetry, theatre, and literature are examined from this perspective in Turner (1996), Freeman (1997), Dancygier (2005), Sweetser (2006), Cook (2006), Fauconnier (2003), Hiraga (2005), Oakley (1998). Courtroom practice is analyzed in Pascual (2002), religion and magic in Sorensen (1999, 2007), Sweetser (2000). Time is re-examined in Fauconnier and Turner (2008).

11. This is one important reason why notions like metaphor, metonymy, analogy, and counterfactual, applied to surface products, elude rigorous definition. Real data does not fit neatly into such categories, and

In order to construct and use generalized integration networks as humans do constantly in everyday life and also in the creative breakthroughs of art, literature and science, they need at a minimum the following capacities, not attested to date in other species: (1) the cognitive capacity for conceptual integration (double-scope, mirror, simplex, ...); (2) the cultural elaboration over cultural time of entrenched networks (e.g. elaborate metaphorical networks developed for time, anger, death, mathematics, event structure, elaborate grammatical construction networks); (3) templates for specific types of integration, such as the counterfactuals in zoloft networks; (4) the capacity to transmit and evolve entrenched networks and templates over generations through learning; (5) the capacity to conceive and build material anchors that stabilize networks and enable or facilitate their transfer and diffusion.

References

- Alac, M. 2006. How Brain Images Reveal Cognition: An Ethnographic Study of Meaning-Making in Brain Mapping Practice. PhD Dissertation, University of California, San Diego.
- Cook, A. 2006. Shakespeare, the Illusion of Depth, and the Science of Parts: An Integration of Cognitive Science and Performance Studies. PhD dissertation, University of California, San Diego.
- Coulson, S. 2001. *Semantic Leaps: Frame-shifting and Conceptual Blending in Meaning Construction*. Cambridge: Cambridge University Press.
- Coulson, S. & Oakley, T. (eds). 2000. *Conceptual Blending*. Special Issue of *Cognitive Linguistics* 11 (3–4).
- Coulson, S. & Oakley, T. (eds). 2005. Special Issue of *Journal of Pragmatics* 37 (10).
- Dancygier, B. 2005. Blending and narrative viewpoint: Jonathan Raban's travels through mental spaces. *Language and Literature* 14 (2): 99–127.
- Evans, V., Bergen, B. & Zinken, J. (eds). In press. *The Cognitive Linguistics Reader*. London: Equinox Publishing Company.
- Fauconnier, G. 2001. Conceptual integration and analogy. In *The Analogical Mind: Perspectives from cognitive science*, D. Gentner, K. J. Holyoak & B. N. Kokinov (eds), 255–286. Cambridge, MA: MIT Press.
- Fauconnier, G. 2003. Compressions de relations vitales dans les réseaux d'intégration conceptuelle. In *Les Sens et la Mesure*, J.-L. Aroui (ed.), 419–427. Paris: Honoré Champion.
- Fauconnier, G. & Turner, M. 1994. Conceptual projection and middle spaces. University of California, San Diego Department of Cognitive Science Technical Report 9401. [Compressed (Unix) postscript version: 9401.ps.Z].
- Fauconnier, G. & Turner, M. 1998. Conceptual integration networks. *Cognitive Science* 22 (2): 133–187. [updated version in Evans, Bergen & Zinken, 2007].
- Fauconnier, G. & Turner, M. 2002. *The Way We Think*. New York: Basic Books.
- Fauconnier, G. & Turner, M. 2008. Rethinking metaphor. In *Cambridge Handbook of Metaphor and Thought*, R. Gibbs (ed.). Cambridge: Cambridge University Press.
- Feldman, J. 2006. *From Molecule to Metaphor*. Cambridge, MA: MIT Press.
- Freeman, M. 1997. Grounded spaces: Deictic-self anaphors in the poetry of Emily Dickinson. *Language and Literature* 6 (1): 7–28.
- Gentner, D., Holyoak, K. & Kokinov, B. (eds). 2001. *The Analogical Mind: Perspectives from Cognitive Science*. Cambridge, MA: MIT Press.

the categories get defined according to the definer's favorite prototypical features and do not provide the theoretical insight needed at a deeper level.

- Hiraga, M. K. 2005. *Metaphor and Iconicity: A Cognitive Approach to Analyzing Text*. New York: Palgrave Macmillan.
- Hougaard, A. & Oakley, T. (eds). In press. *Mental Spaces Approaches to Discourse and Interaction*. Amsterdam: John Benjamins.
- Hutchins, E. 2005. Material anchors for conceptual blends. *Journal of Pragmatics* 37: 1555–1577.
- Imaz, M. & Benyon, D. 2007. *Designing with Blends – Cognitive Foundations of Software Engineering and Human–Computer Interaction Methods*. Cambridge, MA: MIT Press.
- Lakoff, G. 1987. *Women, Fire, and Dangerous Things*. Chicago: University of Chicago Press.
- Lakoff, G. & Núñez, R. 2000. *Where Mathematics Comes From: How the Embodied Mind Brings Mathematics into Being*. New York: Basic Books.
- Liddell, S. 1998. Grounded blends, gestures, and conceptual shifts. *Cognitive Linguistics* 9 (3): 283–314.
- Liddell, S. K. 2003. *Grammar, gesture, and meaning in American Sign Language*. Cambridge: Cambridge University Press.
- Mandelblit, N. 1997. *Grammatical Blending: Creative and Schematic Aspects in Sentence Processing and Translation*. PhD dissertation, University of California, San Diego.
- Núñez, R. 2005. Creating mathematical infinities: Metaphor, blending, and the beauty of transfinite cardinals. *Journal of Pragmatics*. 37: 1717–1741.
- Oakley, T. 1998. Conceptual blending, narrative discourse, and rhetoric. *Cognitive Linguistics* 9: 321–360.
- Pascual, E. 2002. *Imaginary Trialogues: Conceptual Blending and Fictive Interaction in Criminal Courts*. Utrecht: LOT Dissertation Series, 68.
- Robert, A. 1998. Blending in the interpretation of mathematical proofs. In *Discourse and Cognition: Bridging the Gap*, J.-P. Koenig (ed.). Stanford: CSLI publications.
- Sorensen, J. 1999. *Essence, Schema, and Ritual Action*. PhD Dissertation, University of Aarhus.
- Sørensen, J. 2007. *A Cognitive Theory of Magic*. Altamira Press. (Cognitive Science of Religion Series).
- Sweetser, E. 2000. Blended spaces and performativity. *Cognitive Linguistics* 11 (3/4): 305–333.
- Sweetser, E. 2006. Whose rhyme is whose reason? Sound and sense in Cyrano de Bergerac. *Language and Literature* 15 (1): 29–54.
- Turner, M. 1996. *The Literary Mind*. Oxford: Oxford University Press.
- Williams, R. 2005. *Material anchors and conceptual blends in time-telling*. PhD Dissertation, University of California, San Diego.

Genitives and proper names in constructional blends

Barbara Dancygier

1. Frame metonymy and constructions

The genitive (*'s*) form in English has long been seen as semantically puzzling. It plays a special role as the only case in English which is morphologically marked on nouns, and it has long been noticed to display a very broad array of meanings and uses, which are not easily related to one shared semantic source (Nikiforidou 1991; Taylor 1996; Rosenbach 2002). The recent view of the genitive is that it is a means of establishing a reference point (Langacker 1991; Taylor 1996) for the construct represented by the noun being modified. In what follows, I will describe a somewhat more specific use of the genitive, which emerges as the specific contribution of the genitive to two syntactic constructions,¹ both of which can be represented as conceptual integration networks.

My discussion will rely to a large extent on the concept of a frame, introduced in Fillmore (1985) and further developed in Lehrer and Kittay (1992); I also build on the research which uses frames in the analyses of metonymic expressions of various kinds (cf. Koch 1999; Blank 1999).² More specifically, the contribution the genitive makes to constructional meaning will be discussed in terms of three concepts: *frame metonymy*, *constructional compositionality* and *blending*, used, among others, in recent work on conditional constructions in English (Dancygier 1998; Dancygier and Sweetser 2005) and on reflexive constructions in Polish (Dancygier 2005).

The concept of frame metonymy was originally introduced by Fauconnier and Sweetser (1996) to explain examples such as *The ham sandwich wants his check*. While aspects of examples like this can naturally be accounted for through other approaches to metonymy (most will represent the reference to the customer via the food ordered), what frame metonymy is explicitly including in the analysis is the fact that the whole restaurant frame has to be metonymically evoked in order for its specific sub-frames to be linked.

1. A broader discussion of the concept of a grammatical construction is beyond the scope of this paper; I will thus rely on the formulations proposed by Fillmore and Kay (1999), Fillmore, Kay and O'Connor (1988), and Goldberg (1995, 1997).

2. Similar examples have been discussed in terms of domains and subdomains (cf. Croft 1993; Ruiz de Mendoza Ibañez 2000).

In existing literature the two concepts (*frame* and *metonymy*) are not used consistently. On the one hand, Geeraerts (2002) treats any instance of a frame metonymically, seeing the difference primarily in defining metonymy as a semasiological phenomenon, and frame semantics as a sub-field of onomasiology. Geeraerts bases his discussion on contiguity as the primary feature of metonymy (see also Peirsman and Geeraerts 2006) while the treatment proposed by most other metonymy researchers relies in one way or another on the concept of a domain (very close to a common understanding of a frame). For example, while Croft's account (cf. 1993, 2006) relies on *domain highlighting* (rather than a specific reference shift), the treatment proposed in Ruiz de Mendoza Ibañez (2000) and Ruiz de Mendoza Ibañez and Diez Velasco (2002) talks about domain expansion and domain reduction. However, the understanding of frame metonymy advocated here relies heavily on a much broader understanding of what might constitute a frame – from lexical phenomena discussed in the research mentioned above, through constructional frames as discussed by Dancygier and Sweetser (2005), to contextual frames needed in the analysis of proper names.

In Dancygier and Sweetser (2005) the concept of frame metonymy is applied to constructions, not only to lexical meanings. For example, the pattern of verb forms found in predictive futurate conditionals and temporal adverbial constructions, with the present simple in the protasis and the modal *will* in the apodosis, has become so entrenched in the grammar of English, that it can frame-metonymically signal a predictive reasoning in constructions which lack overt exponents of conditionality or temporal sequence, such as the coordinate or conjunctionless constructions. Consequently, sentences such as *He misses another deadline and he'll be fired* and *He misses another deadline, he'll be fired*, both represent the same predictive reasoning as the constructions with adverbial clauses *If/When he misses another deadline, he'll be fired*. The pattern of verb forms thus contributes a part of the interpretation of a coordinate construction, along with other aspects of form. In fact, as Dancygier and Sweetser (2005) show, the combination of parallel constructs such as NPs can lead to similar reasonings (as in *Another day, another dollar.*). It appears, then, that specific constructional features, such as the presence of a specific conjunction, the verb form pattern, or the combination of parallel syntactic units, can compositionally contribute to the resulting meaning of a variety of different constructions.

Furthermore, constructions such as conditionals can be talked about as clusters of various constructional features, with meaning varying with respect to which of the features are present in any given sentence. While particular constructional features frame-metonymically call up their specific meaning contributions as well as broader construction types they most saliently participate in, the overall meaning of a given sentence cannot be fully determined on the basis of its composition, as other (mainly lexical and contextual) factors influence the final interpretation. Consequently, the emergence of the meaning of a specific construction or expression has to be conceived of as a complex operation of conceptual integration (blending), where the contribution of individual lexical and formal features is used in the way appropriate to the combination present.³

3. An example of the blending analysis of the resultative and caused motion constructions can be found in Fauconnier and Turner (1996).

As Broccias observes in his recent review of cognitive approaches to grammar (2006), blending analysis of constructions (as an alternative to Construction Grammar or Radical Construction Grammar) has the advantage of explaining the integrated emergent constructional meaning while also maintaining constructional access to the more basic, unintegrated levels of conceptualization. (For example, the blending analysis gives a more explicit account of the source of causation in the construal of sentences such as *I boiled the pan dry*, cf. Fauconnier and Turner 1996.) This is just the advantage that Dancygier and Sweetser's analysis builds on. In a sentence such as *He misses another deadline, he'll be fired*, predictive reasoning emerges from a combination of syntactic, morphological, and lexical prompts (iconic clause sequence, verb forms pattern, and a scalar expression *another*), none of which can signal predictiveness on its own. The integrated combination, however, emerges out of the frame-metonymic and/or conceptually basic compositional contributions of the individual expressions.

This paper focuses on two constructions which both rely compositionally on the presence of the genitive form. The first construction, which I will call the *GEN-XYZ construction*, can be exemplified by the sentence *To much of the world, Cambodia has become "Vietnam's Vietnam"*, earlier discussed by Glucksberg (2003). The sentence suggests that the government of Vietnam experienced its military intervention in Cambodia in terms similar to what the US government experienced as a result of its military intervention in Vietnam. Apart from the names of two countries being used to stand for the events that occurred there, the genitive form *Vietnam's* is playing its own role in the interpretation. The other construction, which also relies on a similar use of the genitive, is *One person's X is another person's Y*, represented in *One person's trash is another person's treasure*. It is a fairly common two-clausal construction, typically used to contrast different viewpoints. In what follows, I will discuss the compositional features of both constructions, with specific focus on the use of the genitive and its meaning contribution.

2. XYZ constructions and the genitive

Among syntactic constructions recorded and analyzed in the extant literature, the XYZ construction has received much attention as a result of its interesting blending features and its rich potential for metaphorical usage (cf. Turner 1991, 1998; Fauconnier and Turner 2002). The construction is very commonly used in all varieties of English, which seems to result both from its unusually broad range of meaning applications and from the conceptual salience of the emergent blend. The construction is typically represented by an expression of the form NP_1 is the NP_2 of NP_3 , or *X is the Y of Z* (hence the name, XYZ construction). Example (1) represents the simplest use of the construction:

- (1) Paul is the father of Sally.
 X is the Y of Z
 X (*Paul*), Y (*father*), Z (*Sally*), W (*unmentioned [child]*)
 X is a counterpart of Y; Z is a counterpart of W;
 Y-W relationship is projected into the blend; X-Z relationship thus emerges.

In example (2), the same Y (*father*) is used, but the relationship is not understood in terms of kinship. It is, instead, referring to the aspects of the ‘father/child’ frame which relate to authority, responsibility, love, obedience, etc.

- (2) The Pope is the father of all Catholics.
 X is the Y of Z
 X (*Pope*), Y (*father*), Z (*Catholics*), W (*unmentioned [child]*)
 X is a counterpart of Y; Z is a counterpart of W;
 Y-W relationship is projected into the blend; X-Z relationship thus emerges.

Another example discussed by Fauconnier and Turner (2002) illustrates the observation that the Y-W relationship is not necessarily determined by the lexical items used in the construction, but may be read into it (based on the conceptualizer’s knowledge, associative memory, or creative abilities). Example (3) is the case in point.

- (3) The adjective is the banana peel of the parts of speech.
 X is the Y of Z
 X (*adjective*), Y (*banana peel*), Z (*parts of speech*), W (*unmentioned*)
 X is a counterpart of Y; Z is a counterpart of W;
 Y-W relationship is projected into the blend; X-Z relationship thus emerges.

However, once the relationship is determined, it is projected as the relationship between X and Z. If, for example, the banana peel is construed as something that may cause W to slip (and thus lose stable footing), the adjective might then be seen as the category which makes the whole idea of parts of speech shaky and unstable. Interestingly, the XYZ constructions pose significant questions regarding the interaction of lexical and syntactic forms used. While the ‘X is the Y of Z’ form of the construction prompts for the specific type of blend described above, the role of lexical items (specific realizations of X, Y and Z) is crucial in determining the nature of the Y-W and X-Z relationships. Furthermore, the relationship does not have to be determined by the lexical items as such, but may (as in the case of the banana peel) depend on the frame they will call up in the speaker’s or hearer’s mind. The meanings emerging from the integrated XYZ constructions thus cannot be naturally described as cases of ‘coercion’ (Goldberg 1995, 1997) or ‘override’ (Michaelis 2003), but rely crucially on the frame-metonymic function of the Y noun phrase first, and on the frames prompted for by X and Z. It is the metonymic construal of Y that determines the further construal of W and the relationship between them. Consequently, similarly to the case of the resultative construction mentioned above, XYZ constructions depend in equal measure on the construction’s frame and on the partial, unintegrated frames participating in the blend.

The constructions I discuss below can generally be considered variants of XYZ. I will argue, among others, that such variants, whether with genitives or adjectives, use the frames prompted for in ways which are somewhat different from typical XYZ constructions, but do not affect the nature of the constructional contribution of frame metonymy as such. Furthermore, I suggest that the constructional use of genitives (and adjectives) confirms the efficacy of blending in representing the syntax/lexicon interaction.

While some XYZ constructions seem easy to express with a genitive instead of the *of* prepositional phrase, others are not naturally acceptable in the genitive form. For example, *Paul is Sally's father* seems roughly equivalent to (1), but sentences like ??*The Pope is all Catholics' father* or ??*The adjective is the parts of speech's banana peel* are not readily acceptable and do not immediately strike one as equivalent to (2) and (3). Even if we assume that some part of the acceptability judgment depends on the 'heaviness' of the NPs involved, substituting genitive determiners in their place only helps to a degree (consider ??*The Pope is their father* or ??*The adjective is their banana peel*, even when the context makes it clear who the genitives refer to).

At the same time, other examples, like (4),

- (4) My Viper is my Sharon Stone. It's the sexiest vehicle on the road.⁴

can be used with the genitive, but not with the *of* phrase: ??*The Dodge Viper is the Sharon Stone of George (my brother, all drivers). It's the sexiest vehicle on the road*. All these acceptability contrasts suggest that the genitive is making a meaning contribution to the construction, while also limiting the (otherwise very broad) usability of the construction. At the same time, there is a range of contexts where both constructions are acceptable, though it does not follow that they then mean the same things. It seems that the contrasts between the XYZ and GEN-XYZ are best explained in terms of constructional compositionality and the meaning contribution introduced by the genitive form. The approach should let us reveal the similarities between the two constructions, and thus view their shared formal features in terms of their specific contributions to the meaning as a whole.

At the first blush, GEN-XYZ is indeed very similar to XYZ. As the analysis in (5) suggests, both constructions may be interpreted as relying on the relationship between Y and the unmentioned W to be projected into the blend as the relationship between X and Z:

- (5) Iraq is George Bush's Vietnam.⁵ (??? Iraq is the Vietnam of Bush)
 X is Z's Y
 X (*Iraq*), Y (*Vietnam*), Z (*Bush*), W (*unmentioned [Johnson]*)
 X is a counterpart of Y; Z is a counterpart of W;
 Y-W relationship (*Johnson got into an unpopular war [Vietnam] that he couldn't win or get out of*) is projected into the blend; X-Z relationship thus emerges (*Bush got into an unpopular war [Iraq] that he won't be able to win or get out of*).

The source of the difference thus seems to be related to the nature of the relationship which is projected into the blend. Consequently, what calls for an explanation here is, first, where the relationship to be projected resides, and, secondly, the way in which the genitive affects the understanding of Z's role in the X-Z relationship (and, consequently, W's role in the Y-W relationship).

4. From *Parade* magazine; quoted here after Fauconnier and Turner (2002).

5. Ted Kennedy's (2004) comment, quoted by CNN. All the examples in the remainder of this paper come from various samples of journalistic prose (news commentaries, op-ed pieces, etc.), mostly published in their on-line versions in the last five years.

3. Proper names and experiential viewpoint

The relationship of Johnson to Vietnam (and of Bush to Iraq) has to be constructed on the basis of the conceptualizer's knowledge of the historical facts connected to the US intervention in Vietnam. The word *Vietnam*, then, is used here not to refer to the country as such (whether in terms of geographical location, political system, or social structure), but to the circumstances and results of the military intervention of which it was the target. As I have argued elsewhere (Dancygier 2006) proper names are typically used in this manner – not simply to signal a unique referent, but to activate a unique frame of knowledge associated with the name in question. Under this view, unique reference is a consequence, rather than the core of what proper names are designed to do. For example, the use of the name *Paris*, as in *Paris is an interesting city*, will typically evoke a complex frame, including, among other things, information about its being located in France, being the capital of France, and then, depending on the frame held by a given conceptualizer, a bohemian city, or the home of the Eiffel Tower. But it is also possible, in the accepted usage in American English, to append the name with a further clarification of its location (*Paris, France*, as apposed to *Paris, Texas*, which presumably activates a different 'city' frame), just as it is possible that the hearer will only think of Greek mythology, and associate Paris with the man who abducted Helen of Troy. None of these possibilities changes the fact that in each of its uses *Paris* is a proper name with unique reference, but the choice of the specific reference will depend on the frame activated.⁶ This also explains why proper names are so easily used metonymically – for example, *Paris* could further be used to mean 'the capital of fashion' – since in each case the metonymy activates an aspect of the frame.⁷

One linguistically relevant result of treating proper names as frame-metonymic is the possibility of explaining the use of modifiers with such nouns. Rather than treat the cases of proper names modification as exceptions, we can see these cases as relating to the different ways in which the frame can be used.⁸ Such an understanding of metonymy resembles the treatment proposed by Croft (1993, 2006), but the difference is that the concept of domain highlighting can be fruitfully applied to all cases of metonymy, while the use of proper names calls for a more specific explanation. For example, if *Vietnam* stands for the frame related to the US involvement in the country (much more complex than a standard mapping such as Location for Event might suggest), we can see possible modifiers affecting the metonymic use of the frame, not the specific referent, the country of Vietnam. Indeed, at the various stages of the war in Iraq, political commentators were asking questions such as the ones in (6):

6. Recent work by Marmaridou (1989, 1991, 2002) covers some of these facts in terms of Idealized Cognitive Models.

7. Amusingly enough, in one of Agatha Christie's Poirot stories, *Lord Edgware dies*, the murderer is caught because she fails to recognize the mention of Paris as related to the Helen of Troy frame, which was activated earlier, and starts a discussion of fashion instead. This identifies her as the killer!

8. For a recent discussion of proper names modification, see Vandelanotte and Willemsse (2002), Dancygier (2006).

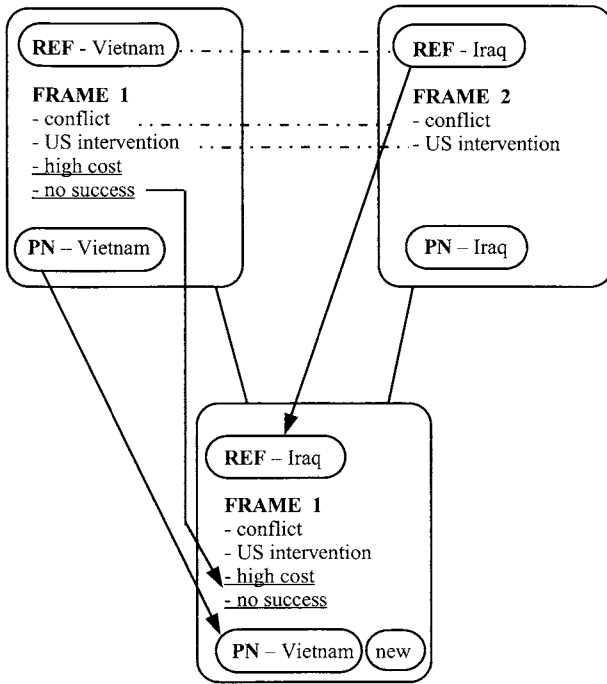


Figure 1. A new Vietnam

(6) Are We Trapped in Another Vietnam? / Is Iraq Becoming a New Vietnam?

which were in fact questions about the new (or another) application of the frame associated with Vietnam, which the public was familiar with. In a nutshell, these were questions about the potential costs and consequences of the US involvement in Iraq. The structure of the blend representing the meaning of the construction is diagrammed in Figure 1.

The facts the figure attempts to describe affect both the form and the meaning of this blend. There are two frames at work signaled by the names of the two countries. They share some structure (e.g. US intervention), but the *Vietnam* frame has structure that the *Iraq* frame does not – the high cost and lack of success. When the name *Vietnam* is used in the construction, it projects its whole frame into the blend, but not the actual referent – the country in Indochina, and not the actual time of the conflict. Iraq (the country) is the referent to which the frame is tentatively applied, but the whole construction talks about Iraq as an emerging frame, possibly modeled after the Vietnam one. The adjective *new* thus refers to the new application of the old frame, now blended with a new situation.

In earlier work (cf. Glucksberg 2003) such a frame-metonymic use of the proper noun *Vietnam* was discussed as the case of metaphor. There is no doubt that it is a use which is different from the (presumably literal) reading, as in *Vietnam has a hot climate*. What seems questionable, though, is the assumption that the use of the word *Vietnam* to describe the story of the conflict is inherently figurative – after all, what is being referred to is a part of what we understand *Vietnam* to be about. It may be the case that a particular language

user does not know enough history to get the whole frame, but that would be true of any such case – after all, Poirot’s success in the story mentioned in footnote 7 lies precisely in the murderer not having the right frame at the right time. In terms of the Lakoffian understanding of conceptual metaphor as a mapping from one domain to another, *Vietnam* is not metaphorical, since the domain remains the same. It is certainly metonymic – frame-metonymic, to be exact, since other aspects of the domain, including the specific referent, are not profiled. As I pointed out above, all proper names seem to represent rich contextual frames, thus selecting a unique referent, but the frame (as opposed to the referent) can be the only part highlighted. Even the simplest reference to a common friend as *George* requires that the frame is shared first, and only then can unique reference be established – otherwise, the interlocutor will inevitably ask *Which George?*⁹

The examples discussed above should have made it clear how the concept of frame metonymy is helpful in understanding the lexical and constructional functioning of proper names. They are different from other nouns in that their conceptualization is less dependent on category structure or the semasiological and onomasiological patterns discussed by Geeraerts (2002). The frames proper names call up rely very heavily on contextual knowledge and the situation of the utterance and may highlight various aspects of the frame in ways which are difficult to predict. While proper names figure prominently in the discussion of metonymic mappings such as Controller for Controlled or Author for Work (as in *Bush invaded Iraq* or *They have a Picasso in their living room*), the explanation of the impact of questions such as *Which George?* is not considered to be part of the equation, and is assumed not to rely on metonymy. However, the frame-metonymic account of such questions will also explain why sentences like *They have a Beethoven in their living room* or *Frodo invaded Mordor* are much less likely to rely on the same metonymic mappings than the Bush and Picasso examples above, even though one can have a Beethoven recording at home and even though Frodo went into Mordor with two ‘combatants’ under his control (Sam and Gollum) and he meant to depose the despot ruling it. All uses of proper names require access to frames, and certain types of metonymies are possible on the basis of those frames, but they rely on more than the simple trigger (such as Controller or Author) to be present in the frame.

The approach recently discussed by Panther (2006) is probably the closest to what I am proposing here. Panther argues for the treatment of metonymic expressions as prompts for indexical relationships. The linguistic form used, says Panther, is a vehicle which, via the meaning of the source, provides access to the target meaning. The target meaning, however, may be dependent on other components of the domain, while the whole process relies heavily on background knowledge and the context of the utterance. Furthermore, the target “conceptually integrates the source meaning, allocating it a backgrounded status” (2006: 153). Panther’s analysis is a very exhaustive and convincing account of how various cases of metonymy can be coherently explained as usage events. Though Panther

9. In the movie *Four Weddings and a Funeral*, the main character, Charles, is seriously scolded by a man to whom he introduces himself. The man is outraged that Charles uses the name which, as the man is sure, refers to someone else. The humor is poignant here: the man seems to assume proper names have unique referents, literally.

does not address the issue of proper names, the description summarized here seems to support my proposal for the treatment of proper names as cases of frame metonymy (with their attendant reliance on background knowledge and the integration of source and target). At the same time, Panther de-focuses the idea of domains, arguing that the boundaries of domains and subdomains are often difficult to defend, and highlighting the role of the conceptualizer in determining what counts as a domain. This is also true about the usage described throughout this paper – though in the cases of proper names the role of the conceptualizer may be even more prominent. I will, however, continue to discuss the examples in terms of frames, in order to demonstrate the specific ways in which different expressions use the conceptual structure prompted for by the proper name as such.

Other examples of proper name modification use the frames in still different ways. For example, denominal adjectives, such as *political*, *emotional*, or *cultural* can be used as in (7):

- (7) More than half the school bond issues in California lose. To win in this state, bonds must earn a “supermajority” – two-thirds of the votes cast. Such a threshold is a political Everest.

In (7), the proper name *Everest* is used to talk not about the mountain in the Himalayas, but about its framing with respect to the domain of climbing. Everest represents the top achievement, and the addition of the adjective *political* suggests that the ‘top achievement’ frame should now be applied to the domain of politics, not climbing. As in the case of a *new Vietnam*, the modifier suggests detaching the frame from its original referent and applying it to a new situation, thus portraying both situations as similar.

Frame-rich nouns seem common in this type of constructional blends. When applied to the domain of politics, nouns such as *Waterloo*, *Rubicon*, or *dinosaur* can all be used in the same construction – the referent is projected from the situation under discussion, while the head noun lexeme is projected from the concept whose frame will be applied to the new referent. Finally, the adjective *political* refers the frame to the domain of politics. Consider the examples in (8) and (9):

- (8) The potential significance of this move, a crossing of the political Rubicon, was evident to Palestinian leaders as well as Israelis.
- (9) Which raises the most interesting and unexpected question about Ted Kennedy: Is he a political dinosaur?

In (8), and in many other similar examples, the noun *Rubicon* is used to call up the frame of ‘a significant change, which creates a new, irreversible state of affairs’ (some users do not seem to even be sure what the original event was and that the Rubicon is a river); the adjective suggests that the ‘significant change’ will occur in the political aspects of the Israeli/Palestinian conflict. In (9), the frame associated with the word *dinosaur*, which could roughly be described as ‘unfit to function well in a new situation’ (that was why dinosaurs became extinct) is again applied to the domain of politics – the expression clearly questions Kennedy’s ability to adapt to the changing political circumstances.

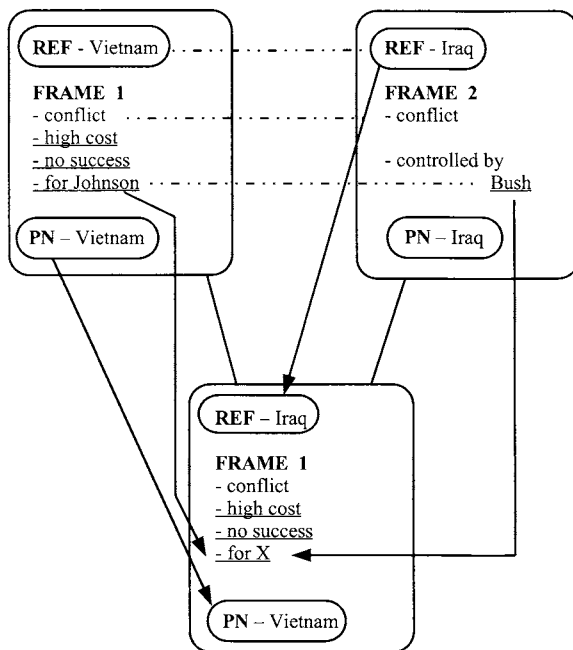


Figure 2. Bush's Vietnam

The constructions described so far (with an instance of *new* or a denominal adjective) are both variants of the XYZ construction.¹⁰ One of the aspects of all these uses is that the frame is prompted by the Y-noun in each case (*father, banana peel, Vietnam, Rubicon, dinosaur*), and is then applied to the new referent – *the Pope, adjective, Iraq, Ted Kennedy*, etc. Furthermore, the constructions which use adjectives like *new* or *political* instead of *of NP* do so also to signal the emerging blend where aspects of the old frame are projected into a new domain. In fact, it is possible to use two kinds of modifiers, with each signaling a different aspect of the blended construction. For instance, it is possible to describe the adoption of a new strategy as *crossing the political Rubicon of the campaign*, to refer to Ted Kennedy as *the political dinosaur of the Democratic Party*, or to describe the Pope as *the spiritual father of all Catholics*. In all of these cases, the referent is talked about in terms of frames primarily attributed to different referents, so the blending processes are similar. It seems, though, that while the *of Z* construction projects a specific relationship into a new situation, the adjectives such as *new* or *political* have a different role – they highlight the assignment of the frame or its crucial parts to a new domain.

The genitive in (5) (*Bush's Vietnam*) is also a modifier which further specifies the way in which the frame associated with *Vietnam* is now applicable to *Iraq* (see Figure 2). It adds a participant to whom the frame is particularly relevant (Bush), and aligns his experience of the frame with the unmentioned participant who is associated with the old frame

10. See Turner (1991, 1998) for a discussion of how the XYZ blend may be expressed with different forms, including noun-noun or adjective-noun combinations.

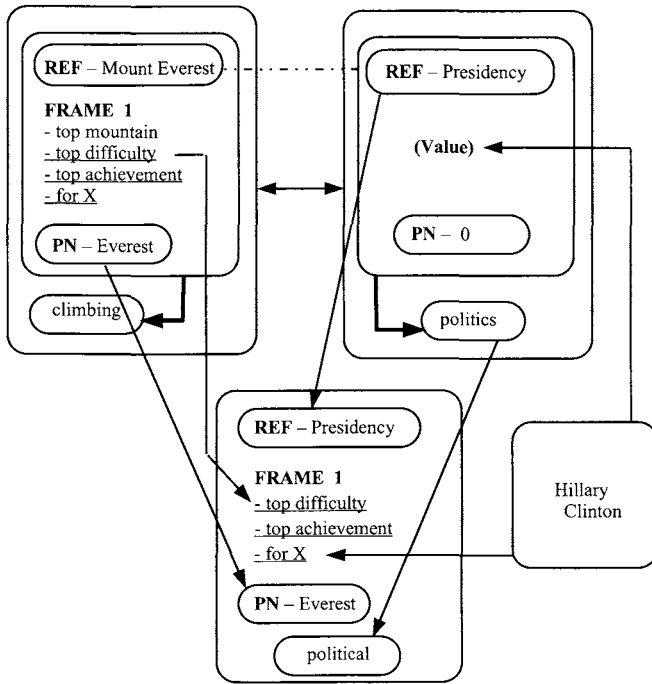


Figure 3. Her political Everest

(Johnson). Also, the same seems to be true of genitive determiners. The determiner *my* in (4) (*my Sharon Stone*) is an example of the same usage, as it identifies the speaker as the person who feels similarly about two unrelated entities – a car and an attractive actress. The meaning that the genitive contributes to the construction seems also quite clear in an example like (10), were it complements the rest of the construction:

- (10) Even though she has gained the prestigious Senate position, it is believed in certain quarters that this is one Hillary who has not yet reached the summit of her political Everest.

The (summit of the) *political Everest* in (10) has the meaning described above – ‘top achievement in politics’. Overall, the fragment quoted here is referring to Hillary Clinton’s ambition to run for the highest office in the US political system. That is, the presidency is presented as ‘the top achievement’ from the point of view of the person represented by the form *her*.

The potential correlation between case and the concept of viewpoint has been explored in detail by Dąbrowska (1997). She argues that a large array of uses of the dative case in Polish is best explained through the concept of the *experiential sphere*. If a parent worries about the possibility of the child becoming sick, she might use an expression *Tylko mi nie choruj* (Just me-DAT not be sick), roughly equivalent to *Just don’t get sick on me*, where the dative structures an understanding of the event as potentially affecting the speaker’s experiential sphere. It seems that the role of the genitive in the English GEN-XYZ constructions is similar – the result is the profiling of a person whose *experiential viewpoint* is accepted in the framing of the new situation in terms of an old one.

Overall, the type of blend suggested in (10) can be represented as in Figure 3. The 'Everest' frame now profiles an additional element – 'for X', where X stands for a potential experiencer of the 'top achievement' frame (the climber getting to the summit of Mount Everest). This frame is blended with the 'presidency' frame, which is contextually given as the one which Senator Clinton would consider the 'top achievement' in her area of activity, which is 'politics'. Hillary Clinton is thus imagined as the future value of the role of 'US President', and occupying that role is seen as the ultimate achievement from her point of view.

In this view, each of the elements of the construction *her political Everest* prompts for its specific element of meaning, which, in turn, contributes to the understanding of the construction as a whole. The word *Everest* frame-metonymically prompts for the frame of 'top achievement' (the proper name stands for its frame, and may imply different referents). The adjective *political* suggests that the achievement in question is in the domain of politics (and not climbing), but it does not specify any situation as suitable. The genitive determiner *her* (contextually referring to Hillary Clinton) specifies the person whose experiential viewpoint presents the situation in question as the 'top achievement'.

As these examples make clear, the syntactic form of the cluster of constructions discussed here, including the GEN-XYZ, is a lot less restricted than other widely discussed constructions, such as the Caused-Motion or The *Way* Construction (Goldberg 1995). What seems particularly variable is the degree to which the referent (X) is syntactically present as the subject of the sentence.¹¹ In an example like (5) *Iraq* is the subject of the sentence, but in (10), the 'presidency' as the target of Hillary Clinton's efforts is only implied, and may in fact not be clear to a reader whose familiarity with Clinton's career is not sufficient. But this seems to be generally true of the situations where frame-metonymy is involved. It is natural to be puzzled by a mention of a person or event which one happens not to be familiar with (for example, many of my own students cannot process new coinages like *Monicagate* or *Bingogate*, because they have not heard of Watergate). The laxness of the 'X is the subject' requirement seems to be the case for all such constructions, including XYZ. A discussion about the Pope may lead to someone saying something like *Well, you would expect the father of all Catholics to get involved*, and the context would provide the X. The X-expression is thus typically the subject of a construction with a linking verb like *be* or *become*, but even if it is not, it has to be possible for a contextually available referent of X to be saliently present (and be focused on) in the preceding discourse.

Furthermore, the constructions crucially depend on the possibility to interpret Y as the carrier of a relational or frame-metonymic concept which can then be used in a projection into another concept. As the wealth of examples given in Turner (1991, 1998) and Fauconnier and Turner (2002) suggest, the blend relies on the listener's/reader's ability to construct the relational similarity which would not exist otherwise. The *banana peel* example quoted above as (3) is a case in point – as Fauconnier and Turner observe, there is no obvious parallelism that would make the blend possible, so it is the reader's task to construct it. But the construction would not work if, let us say, the *of Z* phrase were miss-

11. Thanks to Bill Croft for pointing this out to me.

ing (it would not be acceptable to say **The adjective is the banana peel*), because the blend is prompted by all the lexically present elements of the construction – X, Y, and Z.

The constructions as a group give support to the idea of constructional compositionality, since the different lexical and grammatical ‘ingredients’ (as described above in the analysis of [10]) provide their own meaning contributions to the whole. What seems to be the core of the compositional structure of all the constructions mentioned is the use of Y with a modifier (*of NP, ADJ, or GEN*) which suggests its being understood in a different way. The work done so far on nominal modification as blending (cf. Sweetser 1999; Coulson 2001; Fauconnier and Turner 2002) makes it clear that the interaction between a modifier and a head relies on the relevant concepts connecting only partially and in ways which vary from case to case. The specific nature of the blends involved is what explains why *an apparent mistake* is not really a mistake and how *safe sex* resembles *a safe injection*. However, the modifiers present in the examples of the XYZ family of constructions have a construction-specific role of profiling the domains to which the selected aspects of the Y-frame are to be related. Although the modification integration mechanisms may be similar to those generally available, the modifiers in the XYZ family of constructions have their constructionally determined roles as well. It would not be easily acceptable to simply identify Iraq as Vietnam (as in ??*Iraq is Vietnam*), because there is no identity or equivalence between the meanings of those nouns as such. They can, however, be understood as related if a modifier such as the adjective *new* or a genitive form prompts for the blended concept to be applied to a given domain or viewpoint.

Both the ‘experiential’ use of the genitive in such constructions and the frame-metonymic use of proper names and similar nouns are confirmed in a number of expressions used in news and political commentaries. It is common, for example, for the experiential viewpoint to be assigned to countries, as in *America’s Tsunami* or *Spain’s 9/11*. In both cases the noun calls up a rich event-frame, referring to a relatively recent disaster, and projects the frame onto another event, affecting another country. The December 2004 tsunami in Thailand and Sri Lanka is used as a template frame (massive destruction and death caused by ocean waters entering the shore) to describe the horrors brought to America by the hurricane Katrina (which immediately became a template for the next hurricane, Rita, which was feared to become *another Katrina*, but fortunately proved less powerful). Similarly, the terrorist attack in Madrid was framed, because of its goals and methods, as an event parallel to the attack on the Twin Towers in New York, which came to be termed *9/11*. In both cases the goal of the construction is to frame a recent event in terms of one that the readers are presumably familiar with – even if there are crucial differences. The same construal seems to apply to Glucksberg’s celebrated example *To much of the world, Cambodia has become “Vietnam’s Vietnam”*, which presents the country of Vietnam as invading Cambodia and, ironically, also as subject to experiences analogous to the US experience in the Vietnam conflict. Apart from attesting to the potential of proper names to function frame-metonymically in many ways, the example follows the constructional format of GEN-XYZ.

Naturally, the genitive form can also be used with nouns representing people or groups of people. The defeat at Stalingrad was referred to as *Hitler’s Waterloo*, similar to the above description of the Iraq conflict as *Bush’s Vietnam*. Within Hitler’s experiential sphere the

battle of Stalingrad was indeed comparable to Napoleon's defeat at Waterloo, in that it effectively meant the end of a long war. As in the other cases, the proper name *Waterloo* is used to mean 'irrevocable defeat', which is the most salient aspect of the frame it calls up.

An interesting example to consider is *Engineering's Everest*, the title of an article published in the USC on-line magazine. Without any context, it is hard to guess what is being talked about – one suggestion might be that the text refers to some highly advanced engineering design. However, if this were the case, why would the genitive be used? It would seem more natural to describe it as *the Everest of Engineering*. In fact, the article talks about the efforts of the Dean of the USC School of Engineering to "hoist his school to the very pinnacle of academe – as one of the nation's elite engineering programs". As these words (and the rest of the article) make clear, the word *Engineering* does refer to people and their experience after all, as they work to make their program one of the best (so Everest stands for 'top achievement' again).

To conclude, the contrast between the XYZ construction and the GEN-XYZ relies on the meaning contribution of the genitive, which makes the meaning more specific with respect to the relationship projected. While the GEN-XYZ profiles an experiencing entity, the XYZ may, but does not have to, and if it does, the experiential viewpoint may not be central to the meaning at all. The phrase used by *The New Yorker* journalist to describe Angela Merkel (before her election) was *Germany's Mrs. Thatcher* – apparently in expectation of how she might affect the country's sense of being governed. But it would be equally possible to describe her as *the Mrs. Thatcher of Germany*, if, let us say, her international policies were in focus. The differences may not seem huge in the cases where both forms yield a useful blend, but they are significant. It should be emphasized, however, that both constructions rely compositionally on the frame-metonymic role of the name used as Y, since both constructions expect the Y noun to contribute the frame, not the referent, to the final blend. As a result, a phrase which was commonly used to talk about 'the Iron Lady', *England's Mrs. Thatcher*, is not an example of the GEN-XYZ construction, because the proper name does not prompt for a blend – the frame and the referent are both associated with the same referential expression.¹²

4. XYZ versus GEN-XYZ constructions

The GEN-XYZ construction has been presented above as building on the constructional components present in the XYZ constructions, while also being enriched with the concept of experiential viewpoint, brought into the construction by the use of the genitive. I have also suggested that the XYZ construction differs from GEN-XYZ in that it does not focus on the experiential allocation of the frame, while relying to the same degree on the

12. It is also not necessary that the use of the genitive in this case is contributing an experiential viewpoint. In fact, the examples I have looked at so far seem to suggest that the experiential genitive might be a construction-specific use of the genitive, and thus a special case of the broader 'reference point' definition. Any further investigation of the experiential meaning of the genitive is beyond the scope of the present paper.

frame-metonymic function of Y. In the present section I will consider more examples of the interaction between frame-metonymy and viewpoint.

One of the crucial aspects of the XYZ blends is the contrast between X, which provides the referent, and Y, which prompts for a new frame to be assigned to X. This is further confirmed by examples like (11), where (what looks like) Y does have the function of introducing a new frame, but the referent remains the same. As a result, the examples in (11) are not typical uses of XYZ.

- (11) The first was the Vietnam of the American War, as the Vietnamese call it, ... It was the Vietnam of body counts... .

The fragment comes from a commentary on how the country of Vietnam can mean many different things – or, as I would prefer to describe it, how it can call up different frames. In effect, the sentence talks about “different Vietnams”, while still describing one, generic referent. The *of NP* phrases seem to function similarly to Z phrases elsewhere – they specify the domains to which the framing in question is allocating the particular image of the country of Vietnam. But there is no contrast between X and Y, because the referent of X (Vietnam) is also the referent of Y, even though the *of NP* modifier is suggesting a different, specific sub-framing.

Example (12), for comparison, is a standard case of the XYZ construction, with a human referent of Z. This might suggest that it should be possible to rephrase the sentence with a genitive form, in parallel to the cases like *Germany's Mrs. Thatcher* and *Mrs. Thatcher of Germany*.

- (12) Burton played 136 performances of Hamlet over 18 weeks. The production grossed \$1,250,000 and Elizabeth hailed him the Frank Sinatra of Shakespeare.

In fact, the genitive paraphrase would not be easy to process here. Under one possible interpretation, the genitive form, with its experiential meaning, would suggest that William Shakespeare himself (rather than the domain of theatre history related to the performances of his plays) might be involved. This is not a likely interpretation, given that Shakespeare cannot be presented as even aware of Burton's work without a much more elaborate contextual framing. Even if Shakespeare's name were used to stand for a framing of a different kind, the experiential sense would still be hard to construct, given that the fragment is clearly focusing on the work of Richard Burton.

The contrast between the genitive and the *of*-phrase can be used in many ways. The dubious quality of the prose aside, the author of the text quoted in (13) is playfully exploiting the XYZ construction.

- (13) The Emperor making no response, Ney looked up and observed that his Majesty had fallen asleep. “That settles it,” he sighed. “To-day is the Waterloo of Napoleon Bonaparte. When a man sleeps at a moment like this his friends would better prepare for a wake.” And Ney was right. Waterloo was the Waterloo of Napoleon Bonaparte. The opposing armies met in conflict, and, as the world knows, the star of the great soldier was obscured forever, and France was conquered. (John Kendrick Bangs, *Mr. Bonaparte of Corsica*)

In the fragment, the battle of Waterloo is referred to twice as *the Waterloo of Napoleon Bonaparte* – in the frame-metonymic sense of the ‘final defeat’ discussed above. Interestingly, the text seems to suggest, through the choice of XYZ instead of GEN-XYZ, that the Emperor himself was not aware that the battle would end his campaign, and his career. Ney, as portrayed in the text, has a clear appreciation of the situation, while Napoleon does not. Ney can thus make a statement about the Napoleonic era coming to an end, but Bonaparte himself cannot be presented as experiencing the situation as an inevitable loss. It seems that the choice of XYZ over the GEN-XYZ construction supports the interpretation suggested in the text as a whole in the way consistent with the experiential function of the genitive.

5. One person’s X is another person’s Y

In the discussion of the constructional use of the genitive so far I have made no attempts to relate it to the meaning of the genitive as such. The existing work on the semantics of the genitive suggests that it is difficult to establish an aspect of meaning which would appear in any one of its extremely varied uses. Langacker’s (1991) and Taylor’s (1996) description of the uses of the genitive as reference point constructions provides the broadest possible explanation of the variety of documented meanings. The ‘experiential sphere’ meaning discussed above, which has not been identified as a separate category so far, could easily be seen as a constructionally determined special case of the reference point meaning – indeed, *Bush’s Vietnam* uses Bush’s position in the Iraq conflict as its reference point. Interestingly enough, GEN-XYZ is not the only construction where the genitive form is compositionally important to the overall meaning. The construction which I will refer to as *One person’s X is another person’s Y* is a case in point.

Possibly the most commonly occurring example of the construction is the expression *One person’s trash is another person’s treasure*. It has acquired a usage status very much like that of a proverb, which poses important questions about the source of the generic interpretation. The construction relies on several formal features, such as the use of indefinite noun phrases with *one* or *another* in the genitive form, and on the verb *to be*. The core meaning of the construction is based on two contrasted concepts – X and Y, such as *trash* and *treasure*, which may have little in common in terms of category structure, but represent opposite evaluations. Furthermore, X and Y are not intended to represent two different referents – on the contrary, they are two competing (though not necessarily standard) descriptions of the same referent. The resulting overall meaning of the whole construction is that the referent (whatever it is) could be viewed differently (positively or negatively) by different people. For example, an old book may be dismissed as *trash* based on its worn-out condition, but may be cherished as *treasure* by a book collector or someone who associates fond memories with it. The expressions X and Y are thus used evaluatively, not descriptively, and may be applied to any referent (a book, a piece of clothing, a postcard, but also a friend or a lover, a job offer, an idea, etc.). The use of the genitive highlights the fact that the conflicting evaluations are maintained by different people, with their specific experiential viewpoints. In a sense, the point of the construction is to say that the referent in question may be subject to differing experiential evaluations – positive and negative ones.

One of the important features of the construction is the ability to interpret X and Y as representing evaluative frames – *trash* is a negative term, while *treasure* is positive. However, the terms chosen may have little to do with the actual referent in any specific case. What matters is the consequences of the evaluative judgment – in the case of *trash / treasure*, the point is whether one would want to throw the object away, or keep and cherish it. The examples of opposing pairs I have collected feature, among others: *loss/gain*, *zealot/visionary*, *spending/job*, *drug/poison*, *blather/progress*, *heretic/martyr*, and *propaganda/news*. The range of referents such pairs may describe is very broad, but since the positive/negative framing and its consequences constitute the core of the construction's meaning, the particular choices do not have to strictly correlate with the nature of the referent.

The construction is so deeply entrenched as meaning 'valuable/not valuable to different people' that it is sometimes used jocularly in a way that focuses on the literal descriptive meaning of the terms, instead of their associated frames. Examples (14) through to (17) all build on the negative frames of their X nouns, but instead of the positive term Y (such as *treasure*) they refer to new methods of dealing with X in ways that make it valuable – through recycling or art. This, then, is an example where the X and Y of the construction are treated descriptively, but in reference to the existing evaluative frames. Examples like these seem to confirm the constructional meaning of *One person's X is another's Y*, since the form of the construction itself contributes the evaluative framing.

- (14) One person's trash is another's lesson in recycling.
- (15) One person's trash is another person's ... Seahorse? (*Washington artist recycles scrap metal in his sculptures*)
- (16) One person's garbage is another's power. (*power provided by methane gas from landfills*)
- (17) One person's mess is another's art. (*exhibition of photographs of the mess left behind by hotel guests*)

The evaluative character of the construction is further confirmed by examples like (18) and (19), where X and Y rely on evaluative adjectives like *happy* and *sad*, either used alone or as modifiers. Examples like (18), where X and Y are bare adjectives, also confirm the constructional role of the syntactic form, since outside of the construction a phrase such as **somebody's happy* would not be naturally acceptable. The sentence also confirms the suggested meaning of the genitive here – the only acceptable interpretation is that the adjectives describe the positive or negative feelings attributed to the experiences of different persons.

- (18) One person's happy is another person's sad.
- (19) One person's happy ending can be another's nightmare.

Furthermore, adjectives like *happy* may appear in the construction when the X noun itself is not clearly indicating the kind of evaluation intended. The concept of *anarchy* would often bring negative connotations, but paired with the genitive form and the adjective *happy* (as in [20]) represents the possibility that a person's experience of anarchy may be positive – especially in contrast to the negatively-valued term *fascist*.

- (20) One person's happy anarchy is another person's fascist free-for-all.

To conclude, *One person's X is another person's Y* is a rich construction where, as in other cases, the syntactic form, the grammatical features, and the choice of lexical material all contribute in crucial ways to the construction as a whole. At the same time, the experiential genitive is a particularly interesting aspect of the construction, since it contributes the meaning also found in another construction. It is thus an example of a constructional feature which, while interacting with other features, is also compositionally independent, in that its meaning contribution appears to remain constant across different constructions.

Both constructions discussed above rely on the genitive forms contributing the meaning of 'experiential viewpoint'. As I suggested above, the experiential sense is possibly a special case of what has become known as 'reference point constructions'. It may be easily available outside of the constructions described here, as in expressions such as *John's problem*, when it refers to a situation which John experiences as a problem. However, while outside of the constructions other meanings may be contextually available (*John's problem* may mean a problem he wants to discuss, but not necessarily a problem he is dealing with), the constructions require that the genitive form is referring to the experiential viewpoint. The specific meaning of the genitive is thus constructionally determined, as a result of the specific way in which these constructions rely on frame metonymy and the kinds of blends they prompt for. The genitive is used in just those cases where the frame determines a role of an experiencing or evaluating participant and where that role is further profiled in the blend.

6. Constructions as blends

The constructions described above show important similarities. First, they both rely on the frame-metonymic function of the lexical items to be used. Whether the issue is using a proper name in reference to its contextual frame, rather than to its (unique) referent, or using a common noun to stand for its evaluative framing, both types of constructions require that the lexical items used contribute their framing (and not their typical referents) to the resulting blends. As has been argued in the literature on blending, any language expression is a blend of its form and its content, but the usage described here further suggests that the content represented by a lexical item may be used to foreground selected aspects of its contextual frame rather than its category structure or its referent.

What should also be noticed (though it cannot be further explored here) is that the constructions discussed are (at least at some level) copular constructions, while also being specific varieties of XYZ blends. In fact, such blends seem to rely to a significant degree on the constructional features of copular sentences. The typology of copular constructions discussed in Sakahara (1996) distinguishes three uses: predication, identification, and identity statements. The *One person's ...* constructions are best described as building on the identity statements, but they expand the schema in order to justify the assumed identity of referents of nouns as incompatible as *trash* and *treasure* (*Trash is a treasure* would be difficult to process). The genitives are needed to override (in the sense proposed

by Michaelis 2003) the basic meaning of an identity statement and present the referential identity as independent of the identity as perceived from different viewpoints.

The copular constructions of the kind *Iraq is another Vietnam/Bush's Vietnam* are better described as examples of predications, based on the fact that the subject NP *Iraq* can be understood as referential, while the predicate NPs cannot. The noun *Vietnam* describes *Iraq* in a new way, by giving it a different role, so a copular construction is an appropriate means. The result, however, is that some form of modification is necessary, so that the sentences are prevented from being read as identity statements. Consequently, *??Iraq is Vietnam* is not acceptable, because the new framing or role has to be signaled constructionally, but *Iraq is not Vietnam*, or even *Iraq is no Vietnam* is acceptable, because the framing negated can be retrieved from the context. Furthermore, the use of genitives in such predicational constructions instantiates another interesting case of override – the genitive cannot be understood as signaling a definite NP, because that would violate the copular construction's schema, which explains why *Bush's Vietnam* can be a predicate NP alongside *another Vietnam* or *a new Vietnam*.

The above observations suggest that the constructions discussed here strongly support the conclusion reached in Broccias (2006). Blending is a useful mechanism to account for the emergence of such innocent looking constructions, because it allows us to show how different levels of constructional schematicity yield a coherent and constructionally unique interpretation. GEN-XYZ (with its adjectival variants) and *One person's X* ... all rely on lower level constructions such as genitives, proper names, adjectival modifiers, or determiners, and all participate in overt or backgrounded copular constructions, some of them with negation. Each one of these forms can be subject to its own constructional restrictions, but the pieces of the puzzle may be trimmed or partially redesigned as they make their way through a more complex construction. It does not seem useful to try and rely on some one-way constructional mechanism, such as 'coercion' or 'override', because the specific configuration of constructional features emerges from all the levels of constructional participation, and all the levels, including the final one, are subject to adjustment as new meanings enter the integration network. Perhaps even more importantly, a blending representation allows one to represent the emergent meaning while also maintaining access to the partial constructions. We can imagine a speaker jocularly responding to the *Bush's Vietnam* phrase by saying *So let him keep it. I don't want any part of it*. The use would have to 'undo' the experiential part of the blend, while maintaining the rest of the emergent structure, but it would have to rely on the genitive input to access its other possible interpretations.

The examples discussed here seem to support the usefulness of the concept of constructional compositionality as a correlate of blending. Individual elements of the construction prompt for their meaning contribution, but the structures called up in this way are then subject to standard blending mechanisms – selective projection, emergence of new structure, etc. To conclude, constructional analysis of the kind proposed above is theoretically and descriptively useful in many ways. It leads to a better understanding of the syntactic, lexical and morphological structures involved, and also to a better appreciation of the meaning potential of individual constructions. It may lead to further revisions of well-established grammatical concepts, such as the distinction between proper

and common nouns, the meaning of case, or the role of negation in copular constructions. Most importantly, perhaps, it offers a very specific account of various levels of construction as contributors to the overall meaning.

References

- Blank, A. 1999. Co-presence and succession: A cognitive typology of metonymy. In *Metonymy in Language and Thought*, K.-U. Panther & G. Radden (eds), 169–192. Amsterdam: John Benjamins.
- Broccias, C. 2006. Cognitive approaches to grammar. In *Cognitive Linguistics: Current Applications and Future Perspectives*, G. Kristiansen, M. Achard, R. Dirven, & F. Ruiz de Mendoza Ibañez (eds), 81–118. Berlin: Mouton de Gruyter.
- Coulson, S. 2001. *Semantic Leaps: Frame-Shifting and Conceptual Blending in Meaning Construction*. Cambridge: Cambridge University Press.
- Croft, W. 1993. The role of domains in the interpretation of metaphors and metonymies. *Cognitive Linguistics* 4 (4): 335–370.
- Croft, W. 2006. On explaining metonymy: Comment on Peirsman and Geeraerts. *Cognitive Linguistics* 17 (3): 317–326.
- Dąbrowska, E. 1997. *Cognitive Semantics and the Polish Dative*. Berlin: Mouton de Gruyter.
- Dancygier, B. 1998. *Conditionals and Prediction: Time, Knowledge and Causation in Conditional Constructions*. Cambridge: Cambridge University Press.
- Dancygier, B. 2005. Constructional compositionality and blending: The case of Polish SLVF constructions. *Proceedings of the Thirtieth Annual Meeting of the Berkeley Linguistics Society*: 456–467.
- Dancygier, B. 2006. Proper names in discourse: Modification and constructional blends. Paper presented at the 8th Conference on Conceptual Structure, Discourse and Language. November 2006, San Diego, California.
- Dancygier, B. & Sweetser, E. 2005. *Mental Spaces in Grammar: Conditional Constructions*. Cambridge: Cambridge University Press.
- Fauconnier, G. & Sweetser, E. 1996. Cognitive links and domains: Basic aspects of Mental Space Theory. In *Spaces, Worlds, and Grammar*, G. Fauconnier & E. Sweetser (eds), 1–28. Chicago: University of Chicago Press.
- Fauconnier, G. & Turner, M. 1996. Blending as a central process of grammar. In *Conceptual Structure, Discourse and Language*, A. Goldberg (ed.), 113–130. Stanford: CSLI.
- Fauconnier, G. & Turner, M. 2002. *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. New York: Basic Books.
- Fillmore, C. 1985. Frames and the semantics of understanding. *Quaderni di Semantica* 6: 222–254.
- Fillmore, C. & Kay, P. 1999. Grammatical constructions and linguistic generalizations: The *What's X doing Y?* construction. *Language* 75 (1): 1–33.
- Fillmore, C., Kay, P. & O'Connor, M. C. 1988. Regularity and idiomaticity in grammatical constructions: The case of *Let alone*. *Language* 63 (3): 501–538.
- Geeraerts, D. 2002. Structuring of word meaning I: An overview. In *Lexicology: an International Handbook on the Nature and Structure of Words and Vocabularies*, A. Cruse, F. Hundsnurscher, M. Job & P. R. Lutzeyer (eds), 304–318. Berlin: Mouton de Gruyter.
- Glucksberg, S. 2003. The psycholinguistics of metaphor. *Trends in Cognitive Science* 7 (2): 92–96.
- Goldberg, A. 1995. *Constructions: A Construction Grammar Approach to Argument Structure*. Chicago: University of Chicago Press.
- Goldberg, A. 1997. Relationships between verb and construction. In *Lexical and Syntactical Constructions and the Construction of Meaning*, M. Verspoor, K. D. Lee & E. Sweetser (eds), 383–398. Amsterdam: John Benjamins.

- Koch, P. 1999. Frame and contiguity: On the cognitive bases of metonymy and certain types of word formation. In *Metonymy in Language and Thought*, K.-U. Panther & G. Radden (eds), 139–168. Amsterdam: John Benjamins.
- Langacker, R. W. 1991. *Foundations of Cognitive Grammar. Descriptive Applications*. Volume 2. Stanford: Stanford University Press.
- Lehrer, A. & Feder Kittay, E. 1992. *Frames, Fields, and Contrasts: New Essays in Semantic and Lexical Organization*. Hillsdale, NJ: Laurence Erlbaum Associates.
- Marmaridou, S. S. A. 1989. Proper names in communication. *Journal of Linguistics* 25: 355–372.
- Marmaridou, S. S. A. 1991. What's so proper about proper names? A study in categorization and cognitive semantics. *Parousia Monograph Series* 15. Athens.
- Marmaridou, S. S. A. 2000. *Pragmatic Meaning and Cognition*. Amsterdam: John Benjamins.
- Michaelis, L. 2003. Word meaning, sentence meaning, and syntactic meaning. In *Cognitive Approaches to Lexical Semantics*, H. Cuyckens, R. Dirven & J. Taylor (eds), 163–210. Berlin: Mouton de Gruyter.
- Nikiforidou, K. 1991. The meanings of the genitive: A case study in semantic structure and semantic change. *Cognitive Linguistics* 2 (2): 149–206.
- Panther, K.-U. 2006. Metonymy as a usage event. In *Cognitive Linguistics: Current Applications and Future Perspectives*, G. Kristiansen, M. Achard, R. Dirven & F. Ruiz de Mendoza Ibañez (eds), 147–186. Berlin: Mouton de Gruyter.
- Peirsman, Y. & Geeraerts, D. 2006. Metonymy as a prototypical category. *Cognitive Linguistics* 17 (3): 269–316.
- Rosenbach, A. 2002. *Genitive Variation in English: Conceptual Factors in Synchronic and Diachronic Studies*. Berlin: Mouton de Gruyter.
- Ruiz de Mendoza Ibañez, F. 2000. The role of mappings and domains in understanding metonymy. In *Metaphor and Metonymy at the Crossroads: A Cognitive Perspective*, A. Barcelona (ed.), 109–132. Berlin: Mouton de Gruyter.
- Ruiz de Mendoza Ibañez, F. & Diez Velasco, O. J. 2002. Patterns of conceptual interaction. In *Metaphor and Metonymy in Comparison and Contrast*, R. Dirven & R. Pörings (eds), 489–532. Berlin: Mouton de Gruyter.
- Sakahara, S. 1996. Roles and identificational copular sentences. In *Spaces, Worlds, and Grammar*, G. Fauconnier & E. Sweetser (eds), 262–289. Chicago: University of Chicago Press.
- Sweetser, E. 1999. Compositionality and blending: Semantic composition in a cognitively realistic framework. In *Cognitive Linguistics: Foundations, Scope, and Methodology*, T. Janssen & G. Redeker (eds), 129–162. Berlin: Mouton de Gruyter.
- Taylor, J. R. 1996. *Possessives in English: An Exploration in Cognitive Grammar*. Oxford: Clarendon Press.
- Turner, M. 1991. *Reading Minds: The Study of English in the Age of Cognitive science*. Princeton: Princeton University Press.
- Turner, M. 1998. Figure. In *Figurative Language and Thought*, A. Katz, C. Cacciari, R. W. Gibbs & M. Turner (eds), 44–87. Oxford: Oxford University Press.
- Vandelanotte, L. & Willemse, P. 2002. Restrictive and non-restrictive modification of proprial lemmas. *Word* 53 (10): 9–36.

PART III

Approaches to grammar

Theory and method

What's (in) a construction?

Complete inheritance vs. full-entry models

Arne Zeschel

1. Introduction

In spite of the obvious importance that is accorded to the notion *grammatical construction* in any approach that sees itself as a *construction grammar* (CxG), there is as yet no generally accepted definition of the term across different variants of the framework. In particular, there are different assumptions about which additional requirements a given structure has to meet in order to be recognized as a construction besides being a 'form–meaning pair'. Since the choice of a particular definition will determine the range of both relevant phenomena and concrete observations to be considered in empirical research within the framework, the issue is not just a mere terminological quibble but has important methodological repercussions especially for quantitative research in areas such as corpus linguistics. The present study illustrates some problems in identifying and delimiting such patterns in naturally occurring text and presents arguments for a usage-based interpretation of the term *grammatical construction*.

2. The issue

Different versions of CxG have put forward different elaborations of the generally accepted, yet somewhat unspecific characterization of constructions as 'form–meaning pairs'. Suggestions for a more restricted understanding of the term include the following:

- constructions are *non-predictable* form–meaning pairs (Goldberg 1995: 4; Kay and Fillmore 1999: 4)
- constructions are (fully) *productive* form–meaning pairs (Kay 2002: 3)
- constructions are *entrenched* form–meaning pairs (Croft and Cruse 2004: 288; Langacker 2005: 140; Goldberg 2006: 5; Bybee 2006: 715)
- constructions are *complex* form–meaning pairs (Langacker 1987: 82; Taylor 2002: 561).

In a first attempt at systematizing these proposals, it will be useful to distinguish between those approaches where constructions (in the intended sense) are regarded as the basic

unit of linguistic organization and those which assume that constructions (in the intended sense) are just one element of linguistic knowledge among others. The latter applies for Langacker's Cognitive Grammar which recognizes semantic, phonological and symbolic structures and does not equate *construction* with "symbolic structure" per se – instead, the term is used to denote internally complex (i.e. composite) symbolic structures and is not intended to distinguish elements with unit status from creatively assembled configurations. By contrast, frameworks that go by the name "Construction Grammar" (with capital initials) in the narrower sense commonly assume that language can be accounted for in terms of constructions alone – in the words of Kay and Fillmore (1999: 1), "to adopt a constructional approach is to undertake a commitment in principle to account for the entirety of each language". However, when it comes to the question of which elements will have unit status in such a grammar and why, opinions differ.

One popular answer is the following: any element (and *only* such elements) that cannot be fully reduced to other, more basic elements. As Fillmore, Kay and O'Connor (1988: 502) put it in a foundational paper, "speakers of English have to know what *red* means and that it is an adjective, and they have to know what *ball* means and that it is a noun. They have to know that adjectives can co-occur with nouns in a modification structure (as in a phrase like *red ball*), and they have to know the proper strategies for giving a semantic interpretation to such adjective-noun combinations. But they do not have to know separately, or to be told, what the phrase *red ball* means. That is something which what they already know enables them to find out". On such approaches, *non-predictability* is therefore the defining criterion for constructional status: structures that can be fully reduced to other structures have no independent status but are viewed as productively assembled epiphenomena.

Elaborating on this distinction, Kay (2002) introduces the additional criterion of (full) *productivity* that is intended to separate "true constructions" from "non-productive, non-constructional pattern[s] of coining" (p. 7).¹ The latter term is used for generalizations over groups of stored expressions that are clearly discernible as a pattern and may occasionally license novel formations of the relevant type, but cannot be invoked freely. Kay's example for a pattern of coining is the schema [A as NP] that is implicit in many intensifying expressions of the type *easy as pie*, *happy as a lark*, *dark as night* etc. Kay acknowledges the existence of singleton instances of the pattern that are not plausibly viewed as fixed expressions themselves (cf. his example *wide-eyed as a marigold* from the BNC), but still contends that the pattern itself should not be accorded constructional status since the acceptability of potential instantiating expressions cannot be predicted. This shows that his primary concern is actually with non-predictability, too.

In general, it is probably fair to say that it is mainly computational and/or more formally oriented linguists who see this property as the crucial criterion for constructionhood, and it is typically encountered in discussions of *competence models*. By contrast, construction grammarians who are primarily interested in language as a psychological phenomenon often take a different view. Specifically, proponents of *usage-based models* (Langacker 1990, 2000; see also Croft and Cruse 2004) emphasize that if the goal is to characterize speak-

1. Kay (2002: 2) actually attributes the distinction between "constructions proper" and "patterns of coining" to Fillmore.

ers' linguistic knowledge, the term *construction* should not be restricted to the minimum of strictly non-predictable structures on grounds of aprioristic economy expectations. Instead, speakers are assumed to store large numbers of (more or less) concrete structures in addition to any (putative) general schemas that subsume them (provided these instances are sufficiently *entrenched*, i.e. cognitively routinized). For instance, Bybee (2006:713) observes that “[s]peakers recognize prefabs as familiar, which indicates that these sequences of words must have memory storage despite being largely predictable in form and meaning”. Consequently, Langacker (2005:140) proposes that “an assembly is accepted as part of ‘the grammar’ to the extent that it is psychologically entrenched and conventional in the speech community”. Furthermore, he observes that the elements thus included cannot be neatly partitioned into ‘lexical’ and ‘grammatical’ structures but are more profitably viewed as occupying a particular position on clines in three dimensions, namely degree of generality (schematicity), degree of productivity and degree of compositionality.

The main difference between the two approaches therefore resides in the balance that they strike between aspects of representation and computation: the predictability criterion is typically employed in the context of so-called ‘complete inheritance’ models that seek to formulate maximally parsimonious grammars (at the expense of processing load), whereas the entrenchment criterion is central to so-called ‘full entry’ models that privilege processing economy (through direct retrieval) over storage demands. Metaphorically speaking, both approaches assume that constructions can be characterized as complex network-like structures (inheritance hierarchies or schematized exemplar clusters), but they differ in what is assumed to be ‘in’ these networks, i.e. what constitutes the nodes: in the complete inheritance view, the lattice consists of maximally generalized templates, and there is no redundant representation of specific instances of a given pattern; in the full-entry view, it is the concrete exemplars which are assumed to be stored, whereas more schematic regularities are merely implicit in their instantiations.

It is clear that linguists (especially corpus linguists) who are conducting empirical research within the framework must take sides in this discussion, since the choice of a particular definition will directly influence the obtained results: even though they are sometimes left implicit, any study of course has to formulate sufficiently detailed criteria for including particular observations in the data set, and the formulation of these criteria for the specific purpose at hand in turn reflects general assumptions (also often left implicit) about what distinguishes instances of a construction from certain isomorphic structures in the first place. I will argue that accounts based on non-predictability face a two-fold problem of indeterminacy here when analysing large amounts of noisy naturalistic data: on the one hand, deciding which structures do and which structures don't possess a particular *semantic* feature is often like drawing a line in the sand – there are not always clear-cut tests, and it is often the case that particular aspects will be *more or less salient* in a given instance as compared to other tokens rather than either unambiguously present or absent in a binary fashion. On the other hand, since there is variation between different speakers, there is also variation in the accumulated productions of these speakers/writers that constitute the corpus. As a result, certain properties of the investigated pattern will have the character of statistical tendencies rather than strictly mandatory features. However, adopting the criterion of (non-)predictability forces the analyst to be fully explicit about

precisely which features are required for inclusion and exactly where (i.e. on which level of schematicity) they are encoded, which, as I will show, may be quite difficult to determine. I will illustrate these problems with a corpus study of the fairly inconspicuous expression in bold face in (1):

- (1) *She'll probably also tell us about Tony the Toddler putting Domestos in the jelly, or whatever embarrassing things he did as a kid. Might be **good for a laugh**.* [BNC J1F]

The question to be pursued here is: what is the status of this string?

3. Case study: *good* XP

3.1 A first approximation

At first glance, the expression NP *be good for a laugh* may not seem very interesting – a particular idiom, fully specified except for the subject position, as such quite unremarkable. Like hundreds of other such items, we find it listed in dictionaries such as The Longman Dictionary of Contemporary English (LDOCE), where it has its own entry: “*good for a laugh* – ‘to be enjoyable, amusing’”. If we turn to another dictionary, the Cambridge Advanced Learner’s (CALD), we also stumble across an example sentence containing the phrase *good for a laugh* – here, however, it is listed as an instance of a less specific pattern: “*be good for something* – ‘to be able and willing to provide something’”. If we now return to LDOCE and have a more thorough look at the impressive entry for *good*, a number of further mentions of expressions of the format *good for* NP crop up, among them the following three under the common subentry *good for something*:

- (2) a. ‘able to be used for a particular period of time’
Your passport is good for another three years.
 b. ‘likely to continue living or being useful for a particular time or distance, even though old or not in good condition’
This old truck is good for another 100,000 miles.
 c. ‘likely to give you something or provide something’
Dad should be good for a few bucks. (LDOCE)

Whereas the paraphrase in (2c) is a close variant of the gloss in CALD, both of which are not far away from the more specific meaning of *good for a laugh*, (2a) and (b) are different again, but resemble two further paraphrases that we find in the OED: “capable of producing, valid for etc.”, and “safe to live or last so long, well able to accomplish so much”. The third entry in the OED again resembles (2c): “of a person, that may be relied on to pay so much”.

These meanings are obviously related in some way. As a first approximation, it seems possible to subsume them to the following more schematic characterization:

- (3) NP BE *good for* NP – ‘X can {be used for, produce, provide} Y’

The fact that we are now dealing with a partially schematic structure associated with a number of intuitively related meanings of course raises the question of what it is that we are dealing with in these examples – is it a construction? Is it several constructions? Is it nothing special at all?

For one thing, the fact that we find the pattern listed in several dictionaries seems to suggest that it is somehow more than the sum of its parts, i.e. a string that the lexicographers classified as a meaningful unit that should be part of a dictionary because it must be independently memorized by learners of English. As indicated in the preceding section, the observation that a particular structure possesses an inherent, non-predictable meaning is commonly taken as an indication of constructional status. But are the meanings postulated in (3) really non-predictable? It is widely acknowledged that much of language is considerably vague, so it could be argued that the different context-specific meanings ascribed to the pattern in (3) are simply the product of flexible inferential enrichments applying to vague but otherwise perfectly compositional semantics. An argument against this suggestion is (4):

(4) ?*Might be bad for a laugh.*

If expressions like *be good for* NP were indeed assembled and evaluated fully compositionally, it is difficult to see why *bad for* NP should not work equally well, here giving the meaning that something will probably *fail* to amuse somebody. Note that the problem does not just arise for *good for a laugh*, which belongs to the class of substantive idioms that are known to often disallow lexical substitutions: ?*Dad should be bad for a few bucks* is not a conventional way of saying that *Dad* is 'unable and unwilling to provide something' either. Taking these observations as an indication that we might indeed be dealing with a partially schematic prefab here, we are now faced with the question of what the precise formal and semantic specifications of this structure are. The following section illustrates some problems involved in inducing these properties from corpus data.

3.2 Problems

As it turns out, the semantic characterization proposed in (3) becomes too narrow once we move beyond the dictionary examples reported above and consult a corpus – among others, running a search of the BNC for the string "good for" also produces examples like the following, which rather mean 'X can receive Y':

- (5) a. *Digital workers were always good for a car loan or a mortgage, perceived as being in secure, well-paid jobs.* [BNC K58]
 b. *By dint of a couple of birdies and a fortunate eagle on the long fourteenth hole, Jack was looking good for a share of the prize money.* [BNC CS4]

While it is not impossible in principle for a schematic construction to comprise constructs with converse subsenses (cf. e.g. transfer vs. privative ditransitives), the semantics of the present target pattern becomes substantially more difficult to characterize in view of this discovery: even assuming an already quite coarse-grained specification in terms of

traditional semantic role labels, if we still have to infer contextually whether X is an AGENT, an INSTRUMENT or a RECIPIENT in the encoded scene, the putative construction obviously does not specify a great deal of information by itself.

This underspecification appears to be more drastic once we consider a complementary difficulty that is illustrated by the contrast in (6):

- (6) a. *For now, pen software is good for data collection where users are picking from pre-defined lists, or marking in check boxes.* [BNC FT8]
 b. *Computers are particularly good for finding, storing and retrieving information.* [BNC HXH]
 c. *Portable computers are good to access information while travelling.* [BNC J75]

Here, examples (6b) and (6c) illustrate the opposite case in which we get roughly the same meaning as in (3) (namely, 'X can be used for Y'), yet now an aspect of the form side is different: rather than having *good for* NP, we now get *good* combining with different types of VPs. In fact, the main difference between *good for* NP on the one hand and *good for* V-ing and *good to* V on the other may appear to be that in the case of the latter two, the specific process that X 'can be used for' is explicitly specified by the verb, whereas in *good for* NP it must be inferred. Example (6a) is actually untypical in this respect since an action noun like *data collection* already specifies the relevant process in itself. However, the difference is quite marked in examples like the following:

- (7) a. *They are good for inflammation of the bowel, breathing difficulties, heart problems, high blood pressure, reducing mucus and it also has a calming effect on the body, especially during nausea.* [BNC CGH]
 b. *Incidentally, this cream cheese mixture, spread on little fingers of bread and cooked in just the same way, is extremely good for a cocktail party as a change from those eternal sausages.* [BNC EFU]
 c. *A basic page printer has a rated life; Canon engines are good for 3,000 pages per month, Ricoh's can handle 5,000, and you should expect them to last around three years at the full rated use.* [BNC G00]

We understand these sentences to mean that the subject of (7a) (i.e. *elderflowers*) is good for *curing* inflammation of the bowel, the *cream cheese mixture* in (7b) is good for *serving* at a cocktail party, and *Canon printers* are good for *printing* 3,000 pages per month. An interesting account of the way in which the different implicit predicates are inferred in such examples is offered by Pustejovsky (1995). Drawing on earlier observations by Katz (1964) and Vendler (1967), he notes that the adjective *good* does not denote some invariable quality but merely ranks an entity with respect to a specific scale:

- (8) a. *Mary finally bought a good umbrella.*
 b. *After two weeks on the road, John was looking for a good meal.*
 c. *John is a good teacher.* (Pustejovsky 1995: 43)

Pustejovsky points out that "[t]he conditions which make an umbrella 'good for something' [...] are very different from those which make John a 'good teacher'", suggesting that the selection of an appropriate semantic dimension is driven by the lexical semantics of the *noun* that *good* applies to. He goes on to develop a model in which a particular set

of parameters in the lexical semantics of nouns (relating to how the noun's denotatum came into being, what it consists of, what its canonical purpose or function is and a few more features) may interact with predicates that the noun is in construction with in order to warrant particular inferences during composition. This approach offers a both more general and more flexible account of the semantics of *good for* NP than (3), which comes out as the context-specific spell-out of the schematic meaning 'suited for V-ing NP'. If we accept this (or something similar which also takes the semantics of the subject NP into account) as a promising approach to how the implicit predicate is inferred in examples like (7), and if we furthermore see no principled difference between the type *good for* NP and the types *good for* VP and *good to* VP except that the latter are more explicit in this respect, it would appear that the hypothesized meaning of the target string is in fact not peculiar to the pattern after all. Interestingly, the meaning 'suited for, fitting, appropriate' instead appears to be a meaning of the lexical item *good*, and quite a special one, too: etymological dictionaries list it as the original, i.e. oldest sense of the adjective, which is furthermore characterized as a two-place predicate 'with a purpose-specification in a dependent construction',² i.e. precisely what we find in the above examples.

At this point, then, we have almost come full circle: having suspected that certain expressions involving the string *good for* NP may instantiate a partially schematic idiom that has unit status in its own right, it has now turned out that the precise formal and semantic specifications of this unit are in fact rather difficult to pin down. What is more, there seems to be a way of arriving at the hypothesized constructional meanings proposed in (3) within a particular approach to *lexical* semantics, and here specifically the lexical semantics of *good*. Now, seeing that we are essentially left with a particular (underspecified) reading of the lexical item *good*, is a constructional perspective on *good for* NP obsolete after all?

Having developed these objections at some length, the second part of my paper will be devoted to showing that this is not the case. On the one hand, I will show that *not* all instances of *good* XP behave alike and that the interpretation of *good* in particular is subject to constructional top-down effects. On the other hand, I will argue that a strict dichotomy between lexical and phrasal constructions is to a certain extent misleading anyway, at least if it is taken to imply that lexical and phrasal characterizations must be mutually exclusive and that phrasal constructions should only be posited if all else fails.

3.3 A constructional approach

The observation that there are particular instances of the three realizations of *good* XP that convey similar meanings cannot obscure the fact that the three structures are not generally interchangeable. To begin with, consider the meaning of *good* in (9):

2. "teleologische verwendung mit zweckangabe in abhängiger konstruktion", Grimms Deutsches Wörterbuch (DWB); online: <http://germazope.uni-trier.de/Projects/WBB/woerterbuecher/dwb/wbgui?lemid=GG27807>

- (9) a. *He gets a bit bolshie in the box, and paws at the floor. If he starts being really difficult and kicking at the partitions we sometimes have to put hobbles on him. Fortunately, he is always good to load.* [BNC ASH]
 b. **He is always good for a load.*
 c. **He is always good for loading.*

(9a) is an example of ‘tough movement’ or Object-to-Subject Raising (OSR for short), so called because the classical transformational analysis assumed that the subject of such clauses was derived from an embedded object position in deep structure (i.e. *This book is easy to read* – *It is easy to read this book*). Cognitive Grammar and construction grammar analyses of OSR (Langacker 1995; Hilpert and Koops 2005) have argued that the ‘raised’ variant is an independent and inherently meaningful construction that is typically used to “describe the quality of the experience of the subject in someone who interacts with it in the way specified by the oblique complement” (Langacker 1995:51), and which therefore imposes certain semantic constraints on the slot containing the ‘raised’ predicate. In (9a), the most congruent of the dominant meanings associated with this slot (i.e. ‘easy’ rather than ‘difficult’) is coerced on *good*, indicating that *good* behaves like such different predicates as e.g. *unproblematic*, *a bitch* and *horrible* here in that it can be construed as matching the semantic requirements of this slot. Consequently, (9a) receives its interpretation in virtue of an interaction of the semantics of *good* with the semantics of the OSR-construction, which is why it cannot be paraphrased by either (9b) or (9c). Next, consider the meaning of *good* in (10):

- (10) a. *MPs were usually made to feel welcome; they were always good for a comment (Outrageous, says senior Tory MP), or a piece of gossip.* [BNC HNK]
 b. **They were always good for commenting.*
 c. **They were always good to comment.*

Here, the CALD paraphrase ‘able and willing to provide’ quoted in 3.1 seems quite appropriate, especially if we cut out the ‘provide’ and leave the predicate to be inferred as suggested above (‘make’ in (10)). As illustrated by (10b) and (10c), the meaning ‘able to + PRED’ is not compatible with the other two structures.

On the other hand, the existence of such differences does not preclude the possibility that there may also be points of overlap, as argued for the ‘suited (+ PRED)’-reading in Section 3.2:

- (11) a. *The music was extremely good to dance to and the skins were excellent dancers, although they turned to the slow, deep soul music for close dancing.* [BNC ARP]
 b. *It was good for dancing.*
 c. *It was good for a dance.*

We can therefore hypothesize that on the one hand, the three meanings of *good* observed in (9) to (11) are not associated with any of the three formal patterns directly, but rather arise in the context of different semantic role configurations (notably with different types of subject arguments). On the other hand, each of the three complementation patterns in turn appears to be restricted to a particular subset of such configurations: for instance, *good for NP* seems fine with AGENT subjects (*good for a comment*), but not with PATIENTS

Table 1. Significant combinations of structural pattern and semantic role of subject NP

	Attracted subjects	Repelled subjects
<i>good for</i> NP	AGENT (p < .001)	PATIENT (p < .001)
<i>good for</i> VP	–	AGENT (p < .01)
<i>good to</i> VP	PATIENT (p < .001)	INSTRUMENT (p < .01)

or THEMES (**good for a load*), whereas *good to* VP works well for the latter (*good to load*), but is not compatible with AGENTS (**good to comment*).

In order to investigate this possibility, I conducted a corpus study and extracted all 3566 instances of *good to* (2020 tokens) and *good for* (1546 tokens) from the BNC, by far the most of which were unwanted hits for present purposes. The criterion for including a particular observation in the study was semantic fit with one of the three relevant paraphrases of *good*, i.e. 'able to + PRED', 'suited (+ PRED)' and 'easy'. Most of the examples thus *excluded* involved benefactive uses (12a), constructions with an extraposed expletive subject and the meaning 'desirable' (12b), combinations of these two patterns (12c) and expressions where the XP following *good* was actually licensed by a preceding *too* (12d):

- (12) a. *Maastricht is good for Britain says Major.* [BNC K1Y]
 'beneficent'
- b. *It would be good to meet up sometime.* [BNC GXG]
 'desirable'
- c. *It is not good for a player to be considered a poor sport.* [BNC K5A]
 'beneficent'/'desirable'
- d. *That sounds too good to be true.* [BNC ABJ]
 'desirable'

Apart from that, there were also numerous other idioms such as *hold good for something*, *bode good for something*, *make good for something* etc. which likewise did not convey the requisite meaning. The remaining 373 observations (206 *good for* NP, 106 *good for* VP, 61 *good to* VP) were coded for semantic role of the subject argument. The following categories were used: AGENT/EFFECTOR, INSTRUMENT, THEME, STIMULUS, RECIPIENT, PATIENT, LOCATION, OTHER. The resulting table showed a highly significant interaction between complementation pattern and semantic role of the subject argument ($\chi^2 = 185.44$, $df = 14$, $p < .001$). However, since several of the cells had an expected frequency of less than five, this result should be interpreted with some caution. Still, looking at the individual contributions to this result, there are five individually significant combinations between structural pattern and subject role that stand out from the rest; they are reported in Table 1 (where 'attracted'/'repelled' means 'occurring significantly more/less often than expected').

Taken together with the observations in (9) and (10) that the patterns are not freely interchangeable, these results confirm the suspicion that there are principled semantic grounds for favouring one pattern over the other in a particular context. Specifically, the target pattern *good for* NP is shown to be significantly associated with AGENT/EFFECTOR subjects, even though all other roles that were coded for are in principle possible in this slot too. The most frequent ones are INSTRUMENT (81), AGENT (53) and STIMULUS (25), which make up for 77% of all occurrences (inferred predicate in square brackets):

- (13) a. *So what do you reckon's good for [curing] a hangover, then?* [BNC HW8]
(INSTRUMENT)
- b. *So I doubt if our hero would have been good for [performing] many heroics after swigging that.* [BNC FF0]
(AGENT)
- c. *After all, asking him why my pay's late is always good for [inducing] a laugh.* [BNC JY7]
(STIMULUS)

In fact, expressions with an INSTRUMENT-subject and 'suited'-meanings like (13a) do not seem to be quite the same as (13b) and (13c) – as illustrated in (11), it is *only* the 'suited'-type where there is overlap between *good for* NP and *good for/to* VP, and such expressions also allow the kind of regular lexical substitutions of *good* that are ruled out in (13b) and (13c) (cf. *these pills here are good, but those ones are bad/excellent/lousy for a hangover*).³ It seems reasonable to suspect therefore that the pattern in (13a) is the source structure of the more specialized/ idiosyncratic uses of *good for* NP illustrated in (13b) and (13c).⁴ Having identified expressions like (13) as the dominant subtype of *good for* NP, we can now further zoom in on relevant expressions. For instance, it would also be interesting to see whether individual subtypes have a preference for specific implicit predicates in the oblique phrase. This is what we would expect to find if indeed there is some specialization going on, which is but another way of saying that a previously inferred aspect of relevant expressions comes to be routinely associated with the respective (sub-)pattern and hence semanticized. For reasons of space, I will merely give one example from the AGENT class here, a type that is interesting because agentive subjects are in principle compatible with a large range of predicates. What we find instead is a small number of recurrent scenarios that have been conventionalized for the construction. Consider (14):

- (14) a. *I mean he knew you were good for a few lire more than usual.* [BNC ASN]
- b. *Following a request for a reference in 1989, NatWest wrote to Mr Maitland stating that the company to which he proposed to send the games was good for £5,000 credit.* [BNC AHB]
- c. *Payment is over a number of years. We choose our customers carefully: only those who can underwrite the loan, give pledges that they are good for the money they have borrowed.* [BNC H98]

(14) is about TRANSFER: an AGENT is said to be able to give/supply/restore etc. the oblique THEME argument to a RECIPIENT. The interesting thing about the recurrence of such transfer-implications in agentive *good for* NP-expressions is not so much how they arise (they are invoked by the semantics of the oblique arguments, all of which refer to elements of financial transactions in (14)), but the very fact *that there is a recurrence* of this scenario in the construction – rather than, say, the occurrence of CONSUMING or BREAKING scenarios or whichever other type of event that involves an AGENT. This observation suggests that speakers have quite detailed knowledge about the functions that the individual subtypes

3. I thank Anatol Stefanowitsch for pointing this out to me.

4. Since it was often difficult to sharply distinguish between the putative source meaning and its semantic spin-offs, examples like (13a) were nevertheless included in the data.

of the idiom can be put to in conventional usage. The existence of such restrictions is an argument for a constructional approach, since they are difficult to reconcile with a purely inferential solution according to which more or less anything should be possible here as long as we can recover the implicit predicate from the semantics of the oblique NP.

Another indication that relevant constraints are encoded on the fairly specific level of these individual subtypes is provided by the example that marked the starting point of my investigation, repeated here as (15):

- (15) *She'll probably also tell us about Tony the Toddler putting Domestos in the jelly, or whatever embarrassing things he did as a kid. Might be good for a laugh.*

With 14 out of 208 observations (7%), *laugh* is the most frequently occurring noun in the oblique NP slot; the percentage rises to 11% (22 tokens) when close semantic variants such as *chuckle*, *giggle*, *snigger*, *joke* and *a bit of ribbing* are included as well. In fact, however, such nouns only occur in a particular subtype of *good for* NP-idioms, namely those with STIMULUS subjects. Since there is also an established agentive schema, there should be nothing to prevent us from interpreting (16) as meaning that it is the subject referent who is laughing, rather than being laughed at:

- (16) *It is a world that ought to have vanished, and has largely done so, yet it persists in certain tired imaginations. The forelock-touching peasant is still around in print, and always good for a condescending laugh.* [BNC AHA]

However, this is not what (16) means. If speakers wish to use *laugh* in an agentive setting, they can use the closely related idiom in (17) instead:

- (17) *Droning Dot and nerdish Nigel do have the odd laugh, but there's never been a character who you could call easy going and game for a laugh.* [BNC K37]

This suggests that different subtypes of *good for* NP such as the variants with AGENT and STIMULUS subjects come with different restrictions on the oblique NP-slot – though obviously related, they are not quite the same.

It might be possible to zero in even further on the different usage patterns of such structures, but there is also an end to splitting at some point, at least if we want to be reasonably confident about the general acceptance of the proposed distinctions. In fact, maybe particular aspects of what has been suggested so far are already controversial. Apart from the fact that I am not a native speaker of English and that there is a directly equivalent idiom in my native tongue German that could possibly influence my judgments, this is quite generally what is to be expected at a certain level of detail, and even more so when discussing something as rare and peripheral as the present example: little exposure to a structure means that speakers are not constantly forced to align and possibly accommodate their perceptions of its conventional usage patterns, something that happens automatically for more frequent constructions. As a result, it is well possible that speakers extract slightly different generalizations about the kinds of meanings that can be expressed with this structure, and they will occasionally come across usages that sound deviant and somehow not quite felicitous from their point of view. With the help of corpora, however,

it is possible to identify those patterns in speakers' productions that are *dominant* and, if robust, not compromised by the odd counterexample.

3.4 Implications

Having argued for the general appropriateness of a construction-based approach, it is now time to turn to the predictability vs. entrenchment issue and to consider how the above observations are most plausibly accommodated within an overall model of linguistic knowledge. I will argue that the complete inheritance view with its reliance on the predictability criterion is unsuitable for delimiting the range of elements that speakers really work with in producing and comprehending language (i.e. constructions). By the same token, I will also argue that a certain strategy for modelling the fine-grained aspects of linguistic knowledge that are illustrated by the behaviour of idiomatic chunks such as the present example in formal linguistic complete inheritance models is inappropriate from a cognitive point of view.

Beginning with the first question, it may appear that the above observations do not run counter to a characterization of constructions as form–meaning pairs that are in some respect unpredictable – if indeed there is a highly specific variant with the meaning 'AGENT is able to (+PRED) PATIENT' that behaves differently from a second type with the meaning 'STIMULUS is able to (induce) RESPONSE (in EXPERIENCER)', then that would be a reason for positing two highly specific constructions here that cover relevant expressions. In fact, a schema like 'STIMULUS is able to (induce) RESPONSE (in EXPERIENCER)' is still not specific enough, though: the response is furthermore restricted to certain types of behaviour that imply a particular evaluation of the stimulus on the part of the EXPERIENCER. (In the case of *good for a laugh*, the implication that it is 'enjoyable or amusing'.) This points to a general problem with the non-predictability criterion: if indeed there is an attempt to cover such kinds of phenomena at all, one is forced to make ever more fine-grained subdistinctions in the data, and once these are set up, everything that is sanctioned or licensed by some element of the resulting system should be on equal footing (namely, fine). But this is not the case. Especially in such idiom variants, there are clearly expressions that involve more 'strain' (Langacker 1987: 69f.) in categorization than others because they are further removed from *convention* (i.e. what is *entrenched*). To come back to our example, I believe it is more useful to think of a chunk such as *good for a laugh* as a stored unit that can be subject to certain analogical extensions rather than to postulate a number of fully explicit constructional schemas (or lexical entries – see below) that cover such extensions. As indicated, the BNC contains several substitutions for *laugh* with close semantic variants such as *chuckle*, *giggle*, *snigger*, and web checks with Google quickly produce more far-flung extensions such as *good for a cheer*, *good for a cry*, *good for a puke* etc. The fact that these are much rarer and probably also less acceptable for many speakers directly follows from the fact they are presumably modelled on *good for a laugh* (and more removed from it) rather than 'generated by rule/schema'. Note that this is not to argue against the existence of schematic constructional templates as such: the more variants of this type a speaker encounters, the more likely it is that a schema with some sort of cognitive permanence will be extracted (cf. Langacker 2000: 59f.). The point remains, however, that it is

difficult to account for the connection between acceptability and semantic proximity to an entrenched exemplar in a model that does not recognize such units in the first place.

We have at this point already arrived at the second question, i.e. the implications of such observations for cognitively plausible models of linguistic knowledge. The present paper has argued for an analysis in terms of partially schematic phrasal templates plus fully specific instances, and I have furthermore assumed that such structures are not in principle different from fully schematic constructions. Other approaches do see a principled difference here. For instance, Müller (2006) argues against positing phrasal constructions altogether. Focusing on resultative constructions in German, his argument is that a phrasal solution as proposed by e.g. Goldberg and Jackendoff (2004) interacts with various other phenomena such as constituent reordering and valence-changing processes in undesirable ways because it leads to a strong proliferation of constructional schemas that are needed to license relevant expressions. As an alternative, expanding suggestions by Nunberg, Sag and Wasow (1994) and Erbach and Krenn (1993), Müller makes a case for encoding non-compositional meanings in a special lexical entry of the head of the construction rather than specifying them directly at the phrasal level. Though focusing on a lexically unfilled construction, the paper also touches on more substantive idioms, arguing that “even the tiniest bit of a sentence may be controlled from within a lexical entry” (p. 879): “As it is possible to shift syntactic information around between lexicon and syntactic rules (Constructions), it is also possible to represent semantic information at non-canonical places and, by doing so, to obtain a grammar that can derive the meaning of all utterances compositionally” (p. 877). The approach thus acknowledges the need to account for certain apparently non-compositional phenomena and endorses their treatment in terms of specialized constructions, albeit as ‘lexical constructions’, i.e. elements that are accounted for ‘in the lexicon’. The main motivation behind stipulating the enormous amount of multiple lexical entries that would be needed to cover the full range of idiomatic and collocational idiosyncrasies found in naturally occurring language is thus to maintain a sharp distinction between grammar and lexicon, rule and list. Whereas the assumption of such a design may be desirable in the context of certain contemporary formal approaches like the one in which this particular discussion is framed (HPSG), proponents of Cognitive Grammar and cognitively oriented versions of construction grammar have pointed out that there is no reason to assume that the architectural assumptions and ideals of such models actually mirror the way in which language is instantiated in and processed by the human mind. In other words, having a fully compositional grammar may be desirable from a formal (especially a computational) perspective, but it does not follow that the necessary stipulations carry over to the psychological domain that cognitively oriented linguists seek to describe.

Coming back to the question of lexical vs. phrasal constructions, the classical argument for positing phrasal constructions are constructional coercion effects (Michaelis 2005) and the type of linguistic creativity exhibited by examples like Goldberg’s (1995) famous *He sneezed the napkin off the table*. As I see it, the main problem of a lexical approach here is not so much that positing a special caused-motion entry for a verb like *sneeze* is ‘implausible’, but that this strategy is inevitably *post hoc* and hence not fit to accommodate the inherent flexibility of linguistic categorization. Humans are very adept at establishing

partial correspondences between different elements of conceptual structure (such as the semantics of the verb slot of the caused-motion construction and particular aspects of the conceptual representation associated with the verb *sneeze*) that cannot be exhaustively predicted and enumerated in advance. From a cognitive linguistic perspective, however, Langacker (2005: 147f.) observes that the possibility of construing *sneeze* as matching the semantic requirements of the caused-motion construction is straightforwardly accommodated as a manifestation of the quite general process of conceptual blending (cf. also Fauconnier and Turner 1996). Clearly, though, a verb like *sneeze* can only be ‘made to fit’ into the slot of the construction if the pattern does have independent existence in some way: we can only construct the relevant link and match the (unprofiled) implication of *sneeze* that there is a forceful expulsion of air to the construction’s requirement of there being some sort of force that is causing some sort of motion *if* there is something to map to in the first place. Whether this ‘something’ is then called a construction, a lexical rule or a “defective lexical item” (Jackendoff 2002: 180) is only important insofar as the choice of a particular term will also suggest whether or not the element in question is viewed as something that is fundamentally different from other elements of linguistic knowledge, rather than one endpoint of a continuum. Usage-based construction grammar assumes that there is no such fundamental difference.

4. Conclusion

The present paper has discussed two conflicting views about the crucial criterion for constructional status, i.e. whether constructions are more profitably defined as *non-predictable* or as *entrenched* form–meaning pairs. Departing from a corpus study of a particular schematic idiom, I have argued that speakers have quite detailed perceptions about what can and what cannot be done with such structures, and that such facts are straightforwardly accommodated within an usage-based system (in which entrenchment is the key criterion) but problematic for a maximally lean complete inheritance model that strictly relies on (non-)predictability. What is it that makes item-specific knowledge and stored exemplars relevant? In the case of non-predictable elements, their relevance is obvious: they are what a speaker *must know* in order to speak the language. However, if indeed the goal is to account for *what speakers know*, then non-predictable elements are just the bare minimum, and it is not difficult to show that speakers actually know a lot more: first, I have argued that either fully concrete or only partially schematic ‘fixed expressions’ are the standard of comparison for various analogical extensions that we find evidenced in corpus data (i.e. they are what people seem to work with in production). Second, they are ‘idioms of encoding’ in the sense of Makkai (1972) that speakers ‘must know’ in the same sense that they ‘must know’ opaque idioms of decoding – for instance, idiomatic English has *good for a laugh*, but not **good for laughing*, **good to laugh*, **good with respect to a laugh* or any other conceivable variant that might have become conventionalized instead. Third, research on formulaic language has found that pre-patterned speech, fixed collocations and readymade prefabs also play an important role in discourse/production because they relieve time pressure on the speaker (Wray 2002). Finally, they also impinge

on comprehension, where stored chunks and (possibly conflicting) higher-level schemas compete for activation as the categorizing structure to be selected for a particular target (Zeschel 2008).

Summing up, since all these different aspects point to the significance of concrete exemplars in linguistic knowledge and processing, it would seem strange to exclude them from a model that explicitly seeks to account for “the entirety of each language” (Kay and Fillmore 1999: 1).

References

- Bybee, J. 2005. From usage to grammar: the mind's response to repetition. *Language* 82 (4): 711–733.
- Croft, W. & Cruse, D. A. 2004. *Cognitive Linguistics*. Cambridge: Cambridge University Press.
- Erbach, G. & Krenn, B. 1993. Idioms and support verb constructions in head-driven phrase structure grammar. Claus-Report Nr. 28, Department of Computational Linguistics, Saarland University [also published under Krenn, B. & Erbach, G. 1993, in Nerbonne, Netter, K & Pollard (eds), 365–396. CSLI Lecture Notes Nr. 46. Center for the Study of Language and Information Publications, Stanford, CA.]
- Fauconnier, G. & Turner, M. 1996. Blending as a central process of grammar. In *Conceptual Structure, Discourse, and Language*, A. Goldberg (ed.), 113–130. Stanford: CSLI.
- Goldberg, A. E. 1995. *Constructions. A Construction Grammar Approach to Argument Structure*. Chicago: University of Chicago Press.
- Goldberg, A. E. 2006. *Constructions at Work. The Nature of Generalization in Language*. Oxford: Oxford University Press.
- Goldberg, A. E. & Jackendoff, R. 2004. The English resultative as a family of constructions. *Language* 80 (3): 532–568.
- Hilpert, M. & Koops, C. 2005. A constructional approach to Object-to-Subject Raising. Paper presented at the 9th ICLC, Seoul, Korea.
- Jackendoff, R. 2002. *Foundations of Language*. Oxford: Oxford University Press.
- Kay, Paul. 2002. *Patterns of Coining*. Ms.
- Kay, P. & Fillmore, C. J. 1999. Grammatical constructions and linguistic generalizations: the What's X doing Y? construction. *Language* 75 (1): 1–33.
- Katz, J. J. 1964. Semantic theory and the meaning of 'good'. *Journal of Philosophy* 61: 739–766.
- Langacker, R. W. 1987. *Foundations of Cognitive Grammar. Volume 1: Theoretical Prerequisites*. Stanford: Stanford University Press.
- Langacker, R. W. 1990. *Concept, Image and Symbol. The Cognitive Basis of Grammar*. Berlin: Mouton de Gruyter.
- Langacker, R. W. 2000. A dynamic usage-based model. In *Usage-Based Models of Language*, S. Kemmer & M. Barlow (eds), 1–63. Stanford: Center for the Study of Language and Information.
- Langacker, R. W. 2005. Construction Grammars: cognitive, radical, and less so. In *Cognitive Linguistics. Internal Dynamics and Interdisciplinary Interaction*, F. J. Ruiz de Mendoza Ibanez & M.S. Pena Cervel (eds), 101–159. Berlin: Mouton de Gruyter.
- Makkai, A. 1972. *Idiom structure in English*. The Hague: Mouton.
- Michaelis, L. A. 2005. Entity and event coercion in a symbolic theory of syntax. In *Construction Grammars: Cognitive Grounding and Theoretical Extensions*, J.-O. Östman & M. Fried (eds), 45–88. Amsterdam: John Benjamins.
- Müller, S. 2006. Phrasal or lexical constructions? *Language* 82 (4): 850–883.
- Nunberg, G., Sag, I. & Wasow, T. 1994. Idioms. *Language* 70 (3): 491–538.

- Pustejovsky, J. 1995. *The Generative Lexicon*. Cambridge, MA: MIT Press.
- Taylor, J. R. 2002. *Cognitive Grammar*. Oxford: Oxford University Press.
- Vendler, Z. 1967. *Linguistics in Philosophy*. Ithaca: Cornell University Press.
- Wray, A. 2002. *Formulaic Language and the Lexicon*. Cambridge: Cambridge University Press.
- Zeschel, A. 2008. Lexical chunking effects in syntactic processing. *Cognitive Linguistics* 19 (3): 419–438.

Words as constructions

Ewa Dąbrowska

1. A lexical learnability problem

The average English speaker with secondary school education knows about 60,000 words; many speakers know 100,000 words or more (Miller 1996). 'Knowing a word' involves knowing a variety of things: its phonological form, grammatical properties, meaning, and, for some words at least, the social contexts and genres in which it is normally used (e.g. the word *horsy* is used primarily in informal spoken language, while *equestrian* is much more formal). It is also a matter of degree: a person may have only passive knowledge of a particular word, i.e. be able to recognise it but not produce it, or have only a rough idea of its meaning: for example, one might know that *trudge* is a verb of motion without being aware what specific kind of motion it designates. At the other extreme, many speakers have very detailed representations which enable them to distinguish *trudge* from near-synonyms such as *plod*, *yomp*, and *lumber*.

How is such knowledge acquired? To answer this question, it will be useful to make a distinction between 'basic' and 'non-basic' vocabulary. By 'basic vocabulary' I mean words designating relatively concrete entities which are learned early in development in the context of face-to-face interaction, where the extralinguistic context offers a rich source of information about meaning. In the simplest case, the learner hears a label (*Look! A cat!*) in the presence of a referent (the neighbours' Burmese) and infers that the phonological form [kæt] refers to the animal.¹ Learning relational words such as verbs and prepositions is a more complex process because relations cannot be experienced or conceptualised independently of the entities participating in them (cf. Langacker 1987: 215, 298ff). Moreover, relational words are rarely used in isolation. Thus, learning the meaning of a relational word usually involves performing a sentence-to-world mapping (cf. Gleitman 1990). For example, to learn the meaning of the preposition *on*, the learner must be exposed to sentences such as *The cat sat on the mat* in a context which enables him or her to infer the meaning of the sentence, and to establish correspondences between chunks of phonological structure (e.g. [kæt], [mæt], etc.) and aspects of semantic structure (in this case, the cat and the mat). A further complication arises from the fact that verbs are typically *not* experienced in the presence of the referent: the events described by sentences such as *He*

1. It should be stressed, however, that even such relatively straightforward situations present the learner with many potential difficulties – see Bloom (2000) for an in-depth discussion.

broke it and *Let's go out*, for example, refer to events which occurred either before or after the speech event. However, in all of these cases, learners have access to a variety of situational clues which help them to establish the conventional meanings of the words they are exposed to.

Non-basic vocabulary includes words which are acquired later in development, typically without the benefit of much extralinguistic support. Prime examples of non-basic vocabulary are words for abstract concepts such as *future*, *compute*, *knowledge*, or *aware*, which refer to entities which cannot be directly observed. Another, less obvious, subcategory are words like *scurry*, *ogle*, *capacious*, and *promontory*, which have relatively concrete referents and whose meanings could *in principle* be learned in the same way as basic vocabulary, through exposure during face-to-face interaction with adults in a suitably rich situational context – but which, in practice, cannot be learned in this way because they are simply not encountered in such contexts: words like *scurry* and *capacious* are overwhelmingly used in written texts.

This distinction is, of course, a matter of degree: many words are encountered in written texts as well as in informal interaction; some learners are exposed to richer spoken input than others; and speakers of all ages occasionally encounter new words in face-to-face contexts. The point is that, as their vocabularies increase, language learners have fewer and fewer opportunities for learning words in the context of informal conversation simply because they already know nearly all the words they hear in such contexts (West, Stanovich and Mitchell 1993). Since vocabulary growth does not slow down but actually increases in late childhood and early adolescence (Anglin 1993), it follows that learners must be learning words in non-face-to-face contexts. Hayes and Ahrens (1988) point out that older learners are exposed to new words primarily in written texts: children's books contain 50% more rare words than adult television or the conversation of university-educated adults; and articles in popular magazines contain three times as many rare words as television programmes and adult conversation.

So from about 10 years of age, children encounter most unfamiliar words in written texts and other situations where the amount of extralinguistic information is very limited. This raises obvious learnability issues: how can the learner discover the meanings of words encountered in such contexts? One obvious source of information is explicit definitions: once the learner has become a reasonably competent language user, he or she can learn new words from verbal descriptions provided by other language users. Some words, especially words referring to scientific concepts taught at school, are probably learned in this way; however, it is unlikely that explicit verbal definitions play a very prominent role in lexical development. School-aged children learn 12–15 new words every day (Miller and Gildea 1987; Anglin 1993; Bloom 2000), and we can safely assume that most children are not exposed to anywhere near this number of explicit definitions. Furthermore, most people are not very good at defining words, even words designating relatively concrete concepts. Consider the following definitions produced by five different British undergraduate students:

- (1) a. People do this when they are being big-headed or feeling particularly pleased with themselves.
- b. Move in a dance-like manner.

- c. Jump around in the manner of a loony! To be bouncy, overexcited. Performing reindeer do this.
- d. Walk in an extravagant, showy, arrogant manner, usually in order to attract attention.
- e. Move affectedly. Most often associated with people taking them icky out of ballerinas or camp men. The most common situation would be a camp man trying to get attention.

All of these are definitions of the same lexical item: the English verb *prance*. It is difficult to envisage how a language learner could learn the conventional meaning of the verb from these descriptions (although of course some useful information *can* be gleaned from them).

Definitions found in dictionaries and textbooks are usually more accurate than those produced by ordinary language users, but this doesn't mean that they are always more helpful. For one thing, they often define synonyms in terms of each other. For example, the *Collins English Dictionary* defines *prance* as 'swagger or strut'. If we look up *strut*, we are told that it means 'walk in a pompous manner; swagger', and *swagger* means 'walk or behave in an arrogant manner'. A learner would be able to form a general idea about the meanings of these words from the dictionary – something like 'walk in a pompous or arrogant way' – but not the differences between them. (Note, too, that this definition is not entirely accurate for *prance*, which refers to a walk with exaggerated movements, but does not necessarily imply arrogance: one can prance when one is overexcited or in high spirits.)

Last but not least, children are not very good at learning words from explicit definitions. Consider the following sentences (from Miller and Gildea 1987) produced by children participating in a vocabulary-building programme at school:

- (2) a. I was meticulous about falling off the cliff.
- b. Our family erodes a lot.
- c. Mrs Morrow stimulated the soup.

Miller and Gildea were rather puzzled by such sentences, until they discovered that, according to the dictionary that the children were using, *meticulous* means 'very careful or too particular about small details', *erode* means 'eat out, eat away', and *stimulate*, 'rouse, excite, stir up'. Clearly, the children have not learned the conventional meanings of these words.

How, then, can learners acquire the meanings of non-basic words? There is a growing consensus in the language development literature that non-basic vocabulary is learned through incidental exposure in texts, primarily written texts (Sternberg 1987; Schwanenflugel, Stahl and McFalls 1997; Nagy, Anderson and Herman 1987). The relative success of computational models such as Latent Semantic Analysis (Landauer and Dumais 1997; Landauer 1998) and Hyperspace Analogue to Language (Burgess, Livesay and Lund 1998) demonstrates that such learning is possible, although it is generally agreed that the mathematical algorithms used by the models are unlikely to correspond in any direct way to what the human brain does. We also know that there is a robust correlation between vocabulary size and the amount of reading that a person does (West et al. 1993; Anderson, Wilson and

Fielding 1988) – but, interestingly, not between vocabulary size and the amount of time spent watching television. The most convincing evidence, however, comes from experimental studies demonstrating that performance on vocabulary tests increases if learners are exposed to texts containing words from the test (see, for example, Schwanenflugel et al. 1997; Nagy et al. 1987; Eller, Pappas and Brown 1988; Robbins and Ehri 1994; and Swanborn and de Glopper 1999 for a review).

However, the gains reported in such studies are typically quite small. A meta-analysis of 15 studies of incidental word learning during reading by Swanborn and de Glopper (1999) revealed that the mean probability of a person learning a previously unknown word to a given criterion was 0.15. This figure is probably an overestimate: in many of the studies the participants were given a pre-test assessing their knowledge of the target words before they read the texts containing them, which probably sensitised them to the words, thereby improving learning. The mean learning rate in studies which didn't use a pre-test, or which used a pre-test with distractor items, was 0.11. Furthermore, only one of the studies in the Swanborn and de Glopper sample (Nagy et al. 1987) measured word learning after a week's delay; in all other studies, the vocabulary test was administered immediately after the participants read the passages. Thus, one could argue that these studies measured how good children were at inferring word meaning from context, not how good they were at learning words. In the Nagy et al. study, performance increased by only 5%.

The fact that the increase in knowledge gained from a single exposure in a written text is relatively small is not particularly surprising, given that individual contexts are not very informative (Nagy, Herman and Anderson 1985; Schatz and Baldwin 1986), but performance improves with more exposures (Jenkins, Stein and Wysocki 1984; Robbins and Ehri 1994). Thus, vocabulary learning from context is a slow, incremental process: a learner must encounter a new word in a number of contexts before he or she is able to form a complete lexical entry.

Research on word learning from context suggests that older children and adults are usually better at this than younger children (Swanborn and de Glopper 1999) and that children with larger vocabularies improve more than children with smaller vocabularies (Robbins and Ehri 1994). The properties of the text are relevant, too: for example, learners are more likely to correctly infer the meaning of a particular word if the density of unfamiliar words in the text is low (Swanborn and de Glopper 1999). Finally, high imageability words are learned better than low imageability words, and, interestingly, non-nouns (verbs, adjectives and adverbs) are learned better than nouns (Schwanenflugel et al. 1997). On the other hand, contextual support (how transparent the context is) and text importance (the importance of the sentence containing the word in the story) appear to have no effect on the amount of learning (Schwanenflugel et al. 1997).

What is less clear is exactly how learners construct lexical representations for new words encountered in reading. It is generally agreed that this involves some kind of 'contextual abstraction', but little attempt has been made to isolate the specific clues that learners exploit. Nippold (1998: 18) lists some types of cues that are often available in school textbooks; a selection of items from her list is given in (3) below.

- (3) a. appositives: *Indigo*, a blue dye taken from plants, was sold by Southern plantation owners.
- b. the conjunction *or*: Sir Edmund Hillary climbed to the *summit*, or highest point, of the world's tallest mountain.
- c. metaphor: The bean-shaped *mitochondria* are the cell's power plants.
- d. cause-effect: The pain was *alleviated* as a result of the drugs suggested by the doctor.
- e. participial phrases: The cat, *drenched* by the heavy rain, was distressed.

Note that the cues given in (3a–c) are essentially definitions. Explicit definitions are often available in textbooks, but are not reliably present in other types of texts.² The other cues rely on the learner's ability to make inferences on the basis of real-world knowledge: heavy rain will make a cat wet, drugs can relieve pain, and so on. Being able to make such inferences would allow the learner to formulate a reasonable hypothesis about the meanings of the relevant words. However, Nippold gives no evidence that learners actually use such cues, just notes that they could be used.

Sternberg (1987) does attempt to provide such evidence through two instructional experiments which involved teaching children to attend to specific aspects of context (e.g. temporal, spatial, and causal cues) and to isolate those which are relevant to the meaning of the word. Children who received such training performed better on a subsequent post-test (in which they were required to define new words they encountered in written texts) than a control group who had not. However, it is not clear that the effect was due to attending to the specific clues mentioned by Sternberg – rather than to the fact that the experimental group were encouraged to process the texts more deeply, for example – or how this relates to word learning in the real world, i.e. whether children use the same strategies outside the classroom, and whether the improvement reflects enhanced ability to learn words from context and not simply an enhanced ability to write definitions.

This is not to deny that pragmatic inferencing plays an important role in vocabulary acquisition. The involvement of inferencing processes is largely responsible for the high correlation between vocabulary and IQ,³ and also explains why the ability to learn words from context improves with age. However, there are other sources of contextual information available to the learner which rely on simpler forms of information processing.

First, there is the syntactic frame. Given an unfamiliar word in a sentence with a directional complement (e.g. *He gorp*ed to the park), one can infer that *gorp* probably refers to some kind of motion; the presence of a sentential complement (e.g. *He tam*med that she had left) suggests a verb referring to a mental state or a communication event, and so on. There is considerable evidence that language learners are able to use such cues – indeed,

2. Note, too, that explicit definitions encountered in texts raise similar problems to dictionary definitions.

3. The correlation between scores on the vocabulary subtest of the Wechsler Adult Intelligence Scale and full-scale IQ is .82 (Wechsler 1958:255) – higher than that of any of the other eleven subtests in this battery, and about the same as the correlations between different IQ tests, which average about .77 (Jensen 1998:91). The correlation between scores on Raven's Progressive Matrices, a nonverbal IQ test, and the Peabody Picture Vocabulary Test, is .69 (Jensen 1998:91).

for verbs, the syntactic context is much more informative than the extralinguistic context alone (Gleitman 1990; Gleitman and Gillette 1995; Gillette et al. 1999).

However, the information that syntactic frames provide is very general: it allows learners to identify the broad semantic category of the verb (motion v. transfer v. mental state) but not its precise meaning. Much more specific cues can be gleaned from a word's collocations and semantic preferences, and I would like to suggest that this is the single most important source of information that learners use to learn relational words from linguistic context.

This proposal was inspired by the work of lexicographers such as Sue Atkins (Atkins 1994; Atkins and Levin 1995) who observed that near-synonyms tend to have distinct collocation patterns.⁴ Systematic comparison of these patterns allows lexicographers to bring out the differences in meaning and thus write better definitions; likewise, I suggest, language learners can use the information inherent in typical collocation patterns and semantic preferences to construct lexical representations in their mental lexicons.

To be able to do this, learners and lexicographers alike must first identify typical collocation patterns. This is not a trivial matter, as it involves sifting through vast amounts of information, much of which is irrelevant. Consider the following sentences with the verb *trudge* (all taken from the British National Corpus):

- (4) a. He set out at ten; he viewed as many houses as possible, trudged across miles of fitted carpet and sanded floors, exchanged weary smiles with anxious vendors.
 b. My watch alarm woke us to a finger cold pre-dawn, though I remained only half awake as we trudged through knee-deep snow to the bottom of the Supercouloir, both of us cursing that we had not brought our skis.
 c. Then he and Ranulf trudged wearily off to bed.
 d. Once there, we lifted ourselves and looked at one another, both of us laughing, trudging grass-stained to the top again.
 e. She trudged slowly behind Evelyn, who took the cloth and started to rub out the first word with painstaking precision.
 f. Due to a power blackout, their hotel was in total darkness when they arrived, and they had to trudge up the stairs with their luggage to the 10th floor.

Much of the information in these sentences is irrelevant to determining the meaning of *trudge*. For example, it won't help the learner to know that in the episode described in (4b), the speaker is only half awake, or that the speaker and his companion are cursing that they had not brought their skis; or that in (4d), the walkers were grass-stained and that they were laughing. What is relevant in these sentences is the reference to deep snow in (b), the walkers' weariness in (c), the upwards path in (d) and (f), the slowness of the motion in (e), and the heavy luggage in (f) – but the learner or lexicographer cannot know this until he or she has considered many more sentences.

To assist them in the task of identifying patterns in the data, lexicographers use concordancing programs which pull out corpus sentences containing a particular word and sort them by surrounding context; many such programs also extract collocates and sort

4. For further research exploring the relationship between collocation and meaning, see also Church et al. (1994), Miller and Charles (1991), Divjak and Gries (2006), Gries and Divjak (this volume).

them according to the strength of the relationship with the target word. Language learners, of course, do not have the advantages of modern technology; and moreover, they are presented with exemplars one at a time, which makes the task of comparing them to other exemplars even more difficult.

How then are learners able to isolate typical contexts for a particular word? I suggest that what helps them to accomplish this formidable task is the fallibility of human memory: the fact that we don't normally remember things that we encounter only once or twice (unless they are particularly striking, or highly significant for personal reasons), but we do tend to remember things we are exposed to many times. In other words, memory acts a kind of filter: learners develop robust representations of comparatively frequent collocations like *trudge wearily*, *trudge slowly*, *trudge through the snow* (or, more generally, *trudge through* plus an expression specifying a dense medium such as snow, mud or thick vegetation), *trudge up the stairs* (or, more generally, *trudge UPWARDS*, which is schematic for *up the stairs*, *upstairs*, *up the steps*, *up the hill*, *to the top*); on the other hand, learners do not store rare, perhaps unique combinations such as *trudge across miles of fitted carpet and sanded floors*. The same process allows learners to note that sentences with *trudge* also repeatedly mention the walker wearing heavy footwear, carrying something heavy, covering a considerable distance, and being cold, wet, and miserable.

2. A Cognitive Grammar solution⁵

Thus, the immediate linguistic context contains a wealth of clues about meaning. Critically, much of this information is explicitly mentioned in actual sentences, and thus does not have to be inferred by the learner. Because of this, learning can rely on a relatively simple process of pattern extraction. Clearly, inferencing and real world knowledge also play an important role: a learner who is able to link the information derived from the textual contexts with visual images of people walking through deep snow, or tired or depressed walkers, will have a richer semantic representation of *trudge*; and a learner who is able to glean additional information through inferencing will need fewer exposures to construct an accurate semantic representation. The point is simply that a considerable amount of learning can occur without invoking such computationally demanding processes.

Using distributional cues as described above, a learner would be able to construct a schematic representation such as that depicted in Figure 1b. The figure follows the usual cognitive grammar conventions (cf. Langacker 1987): the boxes represent units; vertical lines represent symbolic relationships; items in capitals represent semantic units; items in phonemic transcription represent phonological units; and '...' represents a maximally schematic phonological unit (a placeholder indicating that some phonological content is present, but not specifying what it is). An additional convention adopted here is the use of

5. The proposal is an application of Langacker's (1987) Cognitive Grammar. It is also broadly compatible with other similar frameworks such as Construction Grammar (Goldberg 1995) and Radical Construction Grammar (Croft 2001). See Langacker (2005) for an in-depth discussion of the similarities and differences between these approaches.

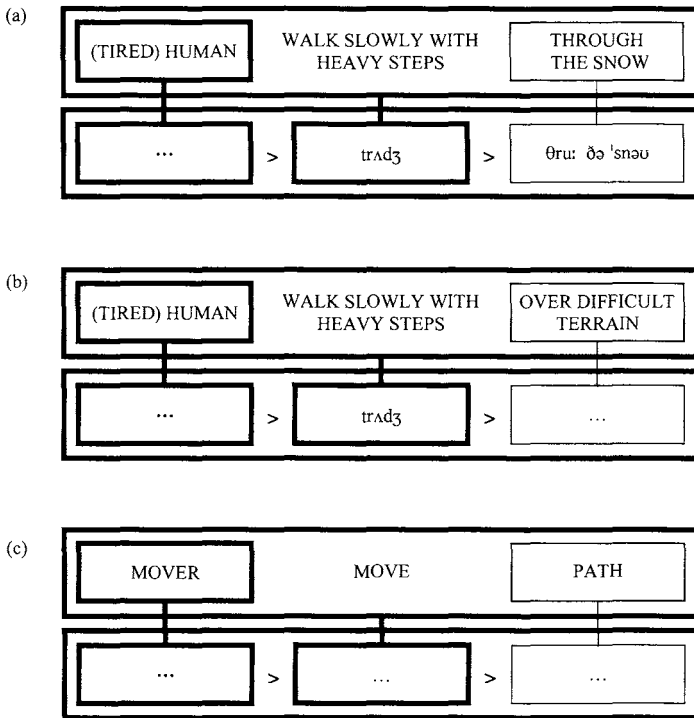


Figure 1. A specific collocation, *trudge through the snow* (a), the lexical representation of the verb *trudge* (b), and the intransitive motion construction (c)

the '<' symbol to represent linear precedence; and for clarity, boxes around symbolic units have been omitted.

The schema in Figure 1b can be regarded as the lexical representation of the verb *trudge*. Such generalized schemas contain representations of the salient participants in the event (in this case, the walker), salient aspects of the setting (difficult terrain), and the phonological form of the linguistic expression used to describe such events. The phonological representation is partially underspecified, in that the segmental content of the phonological subunits corresponding to the walker and the setting is left open; but the unit does specify the ordering of the three subunits. Note that the lexical entry is represented in the same format as constructions and indeed has the same overall structure as the intransitive motion construction (cf. Figure 1c). The only difference between the two representations is that the lexical unit is more specific: it provides more phonological detail and specifies that the mover is human and typically tired, that the motion is slow and bipedal, and happens over difficult terrain. Thus, relational words are, in effect, a special type of construction – one which is partially specified phonologically.

Seeing relational words in this way has several theoretical advantages. Firstly, it makes possible a unified treatment of various aspects of lexical knowledge, including what is traditionally referred to as subcategorization frames and selectional restrictions, as well as frequently co-occurring optional modifiers. All of this information is directly represented

in the schematic specifications of the entities participating in the relationship which are part of the profile of the verb. In this example, the walker is human, and the verb typically, but not always, takes a path expression denoting difficult terrain. The non-obligatory nature of the path expression is represented by thinner lines which indicate that it is less salient than the walker. In addition, specific collocations (e.g. *NP trudge through the snow*, *NP trudge upstairs/up the stairs*) can be represented as independent constructions (cf. Figure 1a) linked to the *trudge* construction via categorizing relationships (Langacker 1987, 2005) or inheritance links (Goldberg 1995), just as *trudge* is linked to the intransitive motion construction. Secondly, seeing relational words as a special type of construction allows a unified treatment of early lexical and grammatical development (acquisition of 'verb islands' and other lexically-specific constructions) and explains the strong correlations between lexical and grammatical knowledge observed in development (e.g. Bates and Goodman 1997): since early constructions are, in effect, big words (cf. Dąbrowska 2000, 2004), we would expect the same mental processes to be involved in their acquisition. Last but not least, as hinted earlier, it explains how, later in development, words can be learned from (written) linguistic context, and allows the analyst to aptly characterize the subtle knowledge that speakers have about the differences between near-synonyms.

On the empirical side, there is a substantial amount of evidence that early in development, children's grammatical knowledge is best characterized as a repertoire of memorised phrases and lexically-specific units such as *CONSUMER-eat-FOOD*, *RUNNER-run-PATH*, *Can I PROCESS?* (Tomasello 1992, 2000, 2003; Lieven, Pine and Baldwin 1997; Dąbrowska 2004). More general constructions such as the transitive, intransitive motion, and Y/N question constructions are acquired later in development by generalizing over the more specific patterns (Tomasello 2000; Dąbrowska 2004).

3. Overview

This paper provides further empirical support for the words-as-constructions view by showing that adult speakers have very specific knowledge about the collocational patterns of particular words which helps them to distinguish between near-synonyms. The specific aspect of linguistic knowledge that will be investigated is verbs of walking or running. English has quite a large number of such verbs, as shown in the list in (5a–b). All of these verbs can be used to describe human bipedal locomotion, although for a few (*gallop*, *trot*, *stampede*, *fly*) this is a secondary sense. There are also a number of more general verbs which are neutral between bipedal and vehicular locomotion (5c), giving a total of about 100 verbs.

- (5) a. walk, **amble**, ambulate, clamber, file, foot it, hike, **hobble**, hoof it, knock about, **limp**, lumber, lurch, **march**, mosey, **pace**, pad, parade, perambulate, **plod**, **prance**, promenade, pussyfoot, ramble, sashay, **saunter**, scuff, **sidle**, shamble, shuffle, skip, skulk, **slink**, slog, **stagger**, stalk, step, **stride**, stroll, **strut**, stump, **swagger**, tiptoe, toddle, traipse, tramp, tread, trek, troop, **trudge**, waddle, yomp

- b. run, beetle, **bolt**, bound, dart, dash, gallop, hotfoot, fly, jog, leg it, lope, romp, rush, scamper, scoot, **scramble**, scud, **scurry**, scuttle, skedaddle, sprint, stampede, trot
- c. move, advance, career, come, decamp, depart, flee, go, hurry, leave, meander, race, roam, rove, skitter, sneak, speed, tear, trek, wander, weave, whisk

Eighteen of these verbs (printed in boldface in the list in (5)) were chosen as the object of the study. All of the verbs are intransitive but typically take directional complements, although most are occasionally used in transitive constructions (as in *the officer plodding the beat*, *posturing crabs who swagger the sea-bed in borrowed shells*, *the designer handbag brigade who strut the Königsallee*, all from the British National Corpus). Apart from *march* all of these are fairly low frequency verbs which are used predominantly in written texts. Adult speakers' knowledge about these verbs was examined by means of a sentence production task (Study 1) and three forced choice tasks (Study 2).

4. Study 1

The first study was an exploratory analysis of speakers' knowledge about the verbs. The 18 verbs were divided into two lists of 9, and 63 undergraduate students (all native speakers of English) were asked to define all the verbs in the set as precisely as they could, and then to use them in sentences illustrating their meaning. One half of the participants were given the verbs from each list. The sentences produced by the participants were collated and coded for characteristics of the walker, path, setting, and manner explicitly mentioned in the sentence. Sentences with non-motion and non-verbal senses of the words (e.g. *I like scrambled eggs*, *I couldn't keep pace with him*) were excluded from the analysis.

In what follows, I report on a subset of the data collected in this way, the illustrative sentences for the nine verbs designating slow movement: *stagger*, *hobble*, *limp*, *trudge*, *plod*, *amble*, *saunter*, *sidle*, and *slink*. Twenty sentences for each verb were included in the analysis. Although this sample is too small to allow firm conclusions to be drawn, it does reveal some suggestive patterns which are summarized in Table 1 and discussed below. For ease of exposition, the nine verbs are grouped into four clusters of nearly synonymous verbs. The division into clusters is based on the author's semantic intuitions and confirmed by an informal similarity judgement study.⁶

6. Ten native speakers were asked to select one or two verbs nearest in meaning to *amble*, *plod*, *sidle*, and *hobble*. At least 8 out of 10 chose *saunter*, *trudge*, *slink*, and *limp*, respectively. The link between *stagger* and *hobble* is weaker, with only two speakers choosing *stagger* as the nearest in meaning to *hobble*. These similarities are also reflected in the pattern of non-target responses observed in Study 2: members of the four pairs of verbs (*amble/saunter*, *plod/trudge*, *sidle/slink*, and *hobble/limp*) were confused with each other much more frequently than with other verbs. The verb most frequently confused with *stagger* was *hobble*; but interestingly, the relationship was asymmetric: that is to say, speakers sometimes supplied *hobble* when the target verb was *stagger*, but never substituted *stagger* for *hobble*.

Table 1. Collocational patterns and semantic preferences in the elicited sentences

	<i>stagger</i>	<i>hobble</i>	<i>limp</i>	<i>trudge</i>	<i>plod</i>	<i>amble</i>	<i>saunter</i>	<i>sidle</i>	<i>slink</i>
Walker									
HUMAN	100	100	95	95	85	100	100	100	65
DRUNK	35	0	0	0	0	0	0	0	0
INJURED/IN PAIN	5	15	40	0	0	0	0	0	0
LEG/FOOT INJURY	0	5	30	0	0	0	0	0	0
CRIMINAL	0	0	0	0	0	0	0	10	15
MALE	65	55	60	40	45	20	60	75	25
OLD	5	50	10	0	0	15	0	0	0
PLURAL/COLLECTIVE	10	0	0	45	25	70	5	5	5
Path									
<i>in/into the room</i>	5	0	0	0	0	0	25	0	0
<i>from/out of the pub/bar</i>	40	5	0	0	0	0	0	0	0
<i>home</i>	40	0	10	20	20	0	0	0	0
<i>off the pitch</i>	0	5	30	0	0	0	0	0	0
<i>along (X)</i>	0	0	5	5	25	35	10	0	5
<i>on</i>	0	0	0	0	30	0	0	0	0
<i>through ... snow</i>	0	0	0	30	0	0	0	0	0
<i>through X</i>	0	0	0	65	5	15	5	0	15
<i>up to PERSON</i>	0	0	0	0	0	0	10	60	0
TOWARDS OPPOSITE SEX	10	0	0	0	0	0	15	5	0
TOWARDS AUTHORITY	0	0	0	0	0	0	0	15	0
TOWARDS	15	25	15	10	10	5	45	80	5
<i>away</i>	0	0	5	0	0	0	0	0	35
AWAY	0	0	5	0	0	0	0	5	45
UPWARDS	5	5	0	10	5	5	0	0	5
no path	5	10	25	0	5	5	0	0	0
Setting									
INDOORS	10	30	20	0	0	5	45	10	5
OUTSIDE	40	40	50	95	70	80	40	5	50
COUNTRY	0	0	0	25	0	75	25	0	5
Manner									
CRUTCHES ETC.	0	15	0	0	0	0	0	0	0

Note: Words in italics correspond to the actual expressions used by the participants; CAPITALS stand for semantic categories. Thus TOWARDS is schematic for *towards*, *to*, *up to*, etc. All the figures given in the table are percentages.

4.1 *Amble* and *saunter*

The dictionary definitions for *amble* and *saunter* are virtually identical: according to the *New Oxford Dictionary of English*, *amble* means 'walk or move at a slow relaxed pace' and *saunter*, 'walk in a slow relaxed manner, without hurry or effort'. However, an examination of the students' sentences reveals some interesting differences. *Amble* is the only verb in the set which is used predominantly with plural or collective subjects, suggesting that this

is an activity one engages in in the company of others; with *saunter*, on the other hand, the subject is virtually never plural. One nearly always ambles outside, typically in the country (*along the trail/footpath, round/across the countryside*); sauntering, in contrast, often occurs indoors. One most often ambles *along*, or *along* something, rarely *away from* or *towards* something (suggesting that one is not going anywhere in particular); but one saunters in a specific direction: *up to* someone (often a person of the opposite sex), *towards* something, or *into a room*. *Amble* is often used with optional modifiers suggesting leisurely activity: *slowly* (2), *without a care in the world* (2), *for an hour, listening to the birds and watching children at play*. *Saunter* also had some modifiers suggesting leisure (*listening to the birds, looking at shop windows*); but there were also modifiers suggesting sexual interest (*sensually*) or a 'studied' casualness (*cool as a cucumber in his new shades, like he had all the time in the world, nonchalantly, unconcerned that he was late yet again*). Last but not least, *amble*, but not *saunter*, appears to be associated with elderly walkers.

4.2 *Plod* and *trudge*

Like *amble*, *trudge* and *plod* are strongly associated with outdoor settings, but unlike *amble*, they tend to be used with modifiers suggesting low energy levels (*wearily, tiredly, after a hard day's work, after a long day at school*). The main difference between the two verbs is in the path: 65% of the sentences with *trudge* described movement *through* something (prototypically *snow*), while the most typical path for *plod* was *along* (with or without a following NP). In addition, *plod*, but not *trudge*, was often used with *on* to indicate continued activity. Another difference is in the choice of subject. All but one of the sentences with *trudge* had human subjects; and interestingly, in the one exceptional sentence, the subject was the coordinate NP *the man and dog*, with a single determiner modifying both nouns, suggesting that they are to be construed as a team. *Plod* seems to allow non-human subjects more freely, especially subjects designating large heavy animals such as elephants and donkeys.

4.3 *Sidle* and *slink*

Both verbs refer to furtive movement, and reflecting this, they were sometimes used with subjects designating criminals (*pickpocket, burglar, robber*) and other disreputable individuals (e.g. *the horny man*). Of all the verbs in this set, *slink* was most frequently used with non-human subjects, typically *a cat*; it is this association which is presumably responsible for the connotations of smooth, gliding movement. With *sidle*, the subject was invariably human. The other significant difference is in the direction of movement. 80% of the sentences with *sidle* describe motion towards something, prototypically *up to* a person of the opposite sex (often with implications of sexual interest), a person in authority or an unsuspecting victim. *Slink*, in contrast, was usually used to describe movement *away* or out of sight (e.g. *into the night*).

4.4 *Hobble, limp and stagger*

These three verbs all refer to an awkward, unsteady movement, but suggest different reasons for the walker's difficulties. In 50% of the elicited sentences with *hobble*, the walker was old (this is a very strong tendency, as the remaining 50% of the subjects were all pronominal); 40% of the sentences with *limp* mentioned some kind of injury, usually to the foot or leg; and 35% of the sentences with *stagger* explicitly stated that the walker was drunk. Some sentences with *hobble* also indicated that the walker used crutches, a Zimmer frame or some other means of support; although references to such aids were not very frequent in absolute terms, they are quite distinctive, since they are not associated with any of the other verbs studied.

Two of the verbs, *stagger* and *limp*, also have strong preferences for particular paths. One typically staggers *from* or *out of* a pub or bar, or *home*: these two paths together account for 80% of the path expressions in the elicited sentences with *trudge* produced by undergraduate students.⁷ For *limp*, the most common path was *off the pitch*;⁸ but the verb was also used fairly frequently without a path expression to describe a manner of walking which is characteristic of a person in the sense that it may be the result of permanent injury.

Thus, while the meanings of these three verbs partially overlap (old people can also limp or stagger, an injured person can hobble or stagger as well as limp, and so on), they have quite distinct prototypical agents: a drunk staggering home after a night out, an injured athlete leaving the game, and an old person unsteady on his/her feet.

4.5 Discussion

The elicited sentences reveal some clear differences in usage patterns which appear to be detailed enough to allow speakers to differentiate between near-synonyms. A relevant question that arises at this juncture is how these patterns compare with those found in 'real' texts. A systematic comparison of the sentences produced by the participants with corpus data is beyond the scope of this paper; suffice it to say that the usage is broadly similar, although the elicited sentences tend to exaggerate patterns found in corpus texts.⁹ For example, in 60% of the elicited sentences with the verb *sidle*, the path was *up to* (a person). *Up to* is also the most frequent collocate of *sidle* in the British National

7. Clearly, this tells us something about the British undergraduate subculture as well as the meaning of *stagger*: one would expect that the results for this verb would be rather different if the participants were old age pensioners.

8. The association of *stagger* with *home* and *from/out of the pub/bar*, and of *limp* with *off the pitch* is very strong, and appears to be giving rise to emergent new senses for these verbs: *stagger* is sometimes used facetiously to refer to going home from a pub even when the walker has not consumed alcohol and is perfectly steady on his/her feet; and *limp* can be used in situations where a player abandons a game because of injury, regardless of whether he or she is actually walking with a limp as they are leaving the pitch.

9. Miller and Charles (1991) observe a similar pattern in their data.

Corpus, but it occurs in only 23% of the corpus sentences. Similarly, *trudge + through ... snow* was attested in 30% of the elicited sentences and only 3% of the sentences in the BNC; for *plod + on*, the relevant figures are 30% and 17% respectively; for *amble + along*, 35% and 15%. These differences are not surprising: participants gave examples of what they considered to be typical usage, while many of the BNC sentences come from literary texts, and hence the language is rather *recherché*. The fact that elicited sentences exaggerate patterns found in the corpus suggests that speakers are aware of what is typical, lending additional support to the idea that lexical representations include knowledge about collocational patterns and semantic preferences.

5. Study 2

The purpose of the second study was to determine how well knowledge of typical collocations predicts performance on other tasks tapping semantic knowledge.

5.1 Method

60 first-year undergraduate students at the University of Sheffield participated in the experiment. All were native speakers of English; none participated in Study 1.

The experiment consisted of three parts: a Definitions task, a Video Clips task, and a Cloze task. The order of the tasks was counterbalanced across participants.

5.1.1 Definitions task

In the Definitions task, participants were given a list of the 18 verbs and their dictionary definitions and asked to choose a verb that went with each definition. For example, for the verb *stride*, participants were presented with one of the following definitions: “walk with long, decisive steps in a specified direction” (*New Oxford Dictionary of English*), “walk with long regular or measured paces, as in haste, etc.” (*Collins English Dictionary and The-saurus*), “walk with long steps, often because one is in a hurry” (*Collins COBUILD English Language Dictionary*, slightly edited), or “walk somewhere quickly with long steps” (*Cambridge International Dictionary of English*). Participants were told that the same verb could be used more than once. There were four versions of the task, each containing definitions from a different dictionary, with the definitions arranged in a different order in each version. Each version was presented to a quarter of the participants. The task took about 5 minutes to complete. One full version of the test is given in Appendix A.

5.1.2 Cloze task

In the Cloze task, participants were presented with 18 sets of five sentences in which the verb was replaced with a blank. They were told that all five sentences in a set contained the same verb, and asked to guess what the verb was; again, the same verb could be used more than once. The 18 verbs were printed at the top of each page. A sample test item is given in (6) below; the complete test can be found in Appendix B. There were four versions of the

test, each containing the same sentence sets in a different order. Each version was given to one-quarter of the participants. The test took about 15 minutes to complete.

- (6) a. I _____ up the stairs.
 b. She _____ through blinding snow.
 c. There was a stream of refugees _____ up the valley towards the border.
 d. He _____ wearily along the path.
 e. We _____ along the muddy track to the top of the hill.¹⁰

The sentences were drawn from examples of usage given in contemporary dictionaries.¹¹ They were thus 'pre-processed', in the sense that they have been selected as typical usages of the verb by the lexicographers who compiled the dictionary; and they are also likely to have been slightly edited. Using such processed examples rather than a random set of sentences from a corpus obviously makes the task of identifying the verb considerably easier for the participants; but note that the purpose of this task was to determine how much participants know about typical collocations, not how good they are at guessing verb meanings using contextual information.

5.1.3 Video clips task

The Video Clips task involved matching the verbs to video clips depicting female actors walking or running in a variety of indoor and outdoor settings (e.g. a car park, a lawn, a formal garden, a large hall, and, for the verb *scramble*, a staircase). Participants were given the following instructions:

You are about to see 18 short 'films', each showing people walking or running in a particular way (*strutting*, *trudging*, *padding*, and so on). Choose the verb from the list below which best describes the way they move and write it in the appropriate blank.

Each 'film' begins with a number and consists of three scenes, each showing the same action. There are short pauses between scenes designed to give you time to think about your answer. Your demonstrator will alert you when the scene begins by saying 'This is 1A' (film 1, scene A), 'This is 1B' (film 1, scene B), and so on.

You can use the same verb more than once. Give only one answer for each film.

Each clip was about 10 seconds long, and there was a 20-second pause at the end of each 'film' during which participants wrote down their answers. The 18 verbs were printed at the top of the answer sheet. The whole test took 18 minutes. All participants completed the same Oxford version of the test.

10. The target verb for this set of sentences is *trudge*.

11. The sentences were taken from the following dictionaries: *Cambridge International Dictionary of English*, *Casell's Modern Guide to Synonyms and Related Words*, *Collins Cobuild English Language Dictionary*, *Collins English Dictionary and Thesaurus* (electronic edition), *The Longman Lexicon of Contemporary English*, *New Oxford Dictionary of English*, *The New Shorter Oxford Dictionary on Historical Principles*, and the *die.net Online Dictionary*.

Table 2. Proportion of target responses in each condition

Verb	Cloze	Definitions	Video clips
march	98	95	98
bolt	93	82	82
pace	92	77	90
stagger	87	77	92
limp	80	70	92
hobble	77	50	83
prance	77	43	78
scramble	73	78	90
scurry	73	80	95
strut	67	60	82
trudge	60	67	62
slink	58	52	70
amble	52	42	43
stride	52	83	78
sidle	52	38	73
plod	42	58	53
saunter	42	33	25
swagger	37	52	72
Mean	67	63	75

5.2 Results and discussion

Table 2 gives information about the proportion of target responses for each verb in each condition. The figures in the table suggest that some verbs (e.g. *scurry* and *scramble*) may be easier to identify on the basis of referential information, while others (e.g. *bolt*) appear to have more distinctive collocates. Overall performance was slightly better on the Video Clips task (75% correct) than on the Definitions and Cloze tasks (63% and 67% respectively). However, such differences are not very informative, since they are to a large extent a direct consequence of the quality of the materials (the use of poor definitions or untypical examples would obviously depress performance on the relevant task) and the intrinsic difficulty of the task (e.g. in the Cloze test, participants had to compare the subjects and path and manner adjuncts in five sentences, which obviously places heavy demands on working memory).

It is much more revealing to compare the correlations between individual participants' scores on the three tasks. As shown in Table 3, performance on the Cloze test was significantly correlated with performance on the other two tasks, but, surprisingly, there is no significant relationship between performance on the Video Clips and Definitions task.¹² In other words, given a person's Cloze score, one can predict their performance on

12. Note that the correlation coefficients are fairly low. This is probably due to the fact that the participants only had partial knowledge of the meanings of the verbs, and therefore had to resort to guessing on some trials; hence, the data are quite noisy. If the test contained more familiar verbs, one would expect higher

Table 3. Correlations between performance on the three tasks

Tasks	Pearson's <i>r</i>	<i>p</i> value
Definitions and Video Clips	0.15	0.243
Cloze and Definitions	0.37	0.005
Cloze and Video Clips	0.37	0.004

the other two tasks; but given the Definitions or Video Clips score, one can only predict the Cloze score. Thus, the results appear to support the hypothesis that knowledge about typical collocations is psychologically more basic.

6. Conclusion

I argued in this paper that relational words such as verbs are constructions, that is to say, units which are complex at both semantic and phonological level. Viewing verbs in this way allows us to give a unified account of how lexical knowledge is acquired and represented, and also helps to explain the otherwise puzzling fact that speakers are able to learn the meanings of new words from purely linguistic contexts. I suggested that they might be able to do this by memorising typical collocation patterns encountered in texts and generalising over them. Previous corpus-based work has shown that sets of near-synonyms have distinct patterns of collocation and colligations (Atkins 1994; Atkins and Levin 1995; Church et al. 1994; Divjak and Gries 2006; Gries and Divjak this volume), and that subjective ratings of semantic similarity are inversely correlated with discriminability of sentential contexts (Miller and Charles 1991). The two experiments described in this paper confirm that speakers have very specific knowledge about the collocations and semantic preferences of individual verbs – even very low frequency verbs which are acquired late in development, which suggests that lexically specific learning continues well into adulthood. Such knowledge appears to be quite subtle, enabling speakers to distinguish between pairs of semantically very similar words such as *amble* and *saunter*, *plod* and *trudge*, *sidle* and *slink*, and *limp* and *hobble*.

Acknowledgements

I would like to thank Jennie Gundill, Kate Moores, and Paula Lago Berdun for acting in the video clips used in Study 2; Marina Samsonova for help with the data collection; and first year students from the Department of English Language and Linguistics at the University of Sheffield for participating in the study. Some of the material reported here was presented at the Ninth International Cognitive Linguistics Conference, the First UK Cognitive Linguistics Conference, and a research seminar at the Department of Modern

overall scores and a significant correlation between performance on the Video Clips and Definitions task; however, the relationship between the Cloze test and the other two tests should still be stronger.

Languages at the University of Trondheim; I would like to thank the audiences in these talks for their helpful comments and suggestions. I am also indebted to Stefan Th. Gries for his comments on an earlier written version of this paper.

References

- Anderson, R. C., Wilson, P. T. & Fielding, L. G. 1988. Growth in reading and how children spend their time outside of school. *Reading Research Quarterly* 23: 285–303.
- Anglin, J. M. 1993. Vocabulary development: A morphological analysis. *Monographs of the Society for Research in Child Development* 58: 1–166.
- Atkins, B. T. S. 1994. Analyzing the verbs of seeing: A frame semantics approach to corpus lexicography. *Proceedings of the Twentieth Annual Meeting of the Berkeley Linguistics Society*: 42–56.
- Atkins, B. T. S. & Levin, B. 1995. Building a corpus: A linguistic and lexicographical look at some near-synonyms. *International Journal of Lexicography* 8: 85–114.
- Bates, E. & Goodman, J. 1997. On the inseparability of grammar and the lexicon: Evidence from acquisition, aphasia, and real-time processing. *Language and Cognitive Processes* 12: 507–584.
- Bloom, P. 2000. *How Children Learn the Meanings of Words*. Cambridge, MA: MIT Press.
- British National Corpus, The*, version 2 (BNC World). 2001. Distributed by Oxford University Computing Services on behalf of the BNC Consortium. URL: <http://www.natcorp.ox.ac.uk/>.
- Brown, L. 1993. *The New Shorter Oxford Dictionary on Historical Principles*. Oxford: Clarendon Press.
- Burgess, C., Livesay, K. & Lund, K. 1998. Explorations in context space: Words, sentences, discourse. *Discourse Processes* 25: 211–257.
- Church, K. W., Gale, W. A., Hanks, P., Hindle, D. & Moon, R. 1994. Lexical substitutability. In *Computational Approaches to the Lexicon*, B. T. S. Atkins & A. Zampolli (eds), 153–177. Oxford: Oxford University Press.
- Collins English Dictionary and Thesaurus* (electronic edition). 2002. London: Collins.
- Croft, W. 2001. *Radical Construction Grammar: Syntactic Theory in Typological Perspective*. Oxford: Oxford University Press.
- Dąbrowska, E. 2000. From formula to schema: The acquisition of English questions. *Cognitive Linguistics* 11: 83–102.
- Dąbrowska, E. 2004. *Language, Mind and Brain. Some Psychological and Neurological Constraints on Theories of Grammar*. Edinburgh: Edinburgh University Press.
- die.net *Online Dictionary*. Available at <http://dict.die.net/>, accessed 10 October 2004.
- Divjak, D. & Gries, S. T. 2006. Ways of trying in Russian: clustering behavioural profiles. *Corpus Linguistics and Linguistic Theory* 2: 23–60.
- Eller, R. G., Pappas, C. C. & Brown, E. 1988. The lexical development of kindergarteners: Learning from written context. *Journal of Reading Behavior* 20: 5–24.
- Gillette, J., Gleitman, H., Gleitman, L. & Lederer, A. 1999. Human simulations of vocabulary learning. *Cognition* 73: 135–176.
- Gleitman, L. R. 1990. The structural sources of verb meanings. *Language Acquisition* 1: 3–55.
- Gleitman, L. R. & Gillette, J. 1995. The role of syntax in verb learning. In *The Handbook of Child Language*, P. Fletcher & B. MacWhinney (eds), 413–427. Oxford: Blackwell.
- Goldberg, A. E. 1995. *Constructions. A Construction Grammar Approach to Argument Structure*. Chicago: The University of Chicago Press.
- Gries, S. T. & Divjak, D. This volume. Behavioral profiles: A corpus-based approach to cognitive semantic analysis. In V. Evans & S. Pourcel (eds). Amsterdam: John Benjamins.
- Hayakawa, S. I. & Fletcher, P. J. 1977. *Casell's Modern Guide to Synonyms and Related Words*. London: Cassell.

- Hayes, D. P. & Ahrens, M. 1988. Vocabulary simplification for children: A special case of 'motherese'? *Journal of Child Language* 15: 395–410.
- Jenkins, J. R., Stein, M. L. & Wysocki, K. 1984. Learning vocabulary through reading. *American Educational Research Journal* 21: 767–787.
- Jensen, A. R. 1998. *The g Factor: The Science of Mental Ability*. Westport, CT: Praeger Publishing.
- Landauer, T. K. 1998. An introduction to Latent Semantic Analysis. *Discourse Processes* 25: 259–284.
- Landauer, T. K. & Dumais, S. T. 1997. A solution to Plato's problem: The Latent Semantic Analysis theory of acquisition, induction and representation of knowledge. *Psychological Review* 104: 211–240.
- Langacker, R. W. 1987. *Foundations of Cognitive Grammar. Volume 1: Theoretical Prerequisites*. Stanford, CA: Stanford University Press.
- Langacker, R. W. 2005. Construction Grammars: Cognitive, radical, and less so. In *Cognitive Linguistics. Internal Dynamics and Interdisciplinary Interaction*, F. Ruiz de Mendoza & S. Peña (eds), 101–159. Berlin: Mouton de Gruyter.
- Lieven, E. V., Pine, J. M. & Baldwin, G. 1997. Lexically-based learning and early grammatical development. *Journal of Child Language* 24: 187–219.
- McArthur, T. 1981. *Longman Lexicon of Contemporary English, The*. Harlow, Essex: Longman.
- Miller, G. A. 1996. *The Science of Words*. New York: Freeman.
- Miller, G. A. & Charles, W. G. 1991. Contextual correlates of semantic similarity. *Language and Cognitive Processes* 6: 1–28.
- Miller, G. A. & Gildea, P. M. 1987. How children learn words. *Scientific American* 257: 86–89.
- Nagy, W. E., Anderson, R. C. & Herman, P. A. 1987. Learning word meanings from context during normal reading. *American Educational Research Journal* 24: 237–270.
- Nagy, W. E., Herman, P. A. & Anderson, R. C. 1985. Learning words from context. *Reading Research Quarterly* 20: 233–253.
- Nippold, M. A. 1998. *Later Language Development: The School-Age and Adolescent Years*. Austin, TX: Pro-ed.
- Pearsal, J. 1998. *New Oxford Dictionary of English*. Oxford: Oxford University Press.
- Procter, P. 1995. *Cambridge International Dictionary of English*. Cambridge: Cambridge University Press.
- Robbins, C. & Ehri, L. C. 1994. Reading storybooks to kindergartners helps them learn new vocabulary words. *Journal of Educational Psychology* 86: 54–64.
- Schatz, E. K. & Baldwin, R. S. 1986. Context clues are unreliable predictors of word meanings. *Reading Research Quarterly* 21: 439–453.
- Schwanenflugel, P. J., Stahl, S. A. & McFalls, E. L. 1997. Partial word knowledge and vocabulary growth during reading comprehension. *Journal of Literacy Research* 29: 531–553.
- Sinclair, J. 1987. *Collins Cobuild English Language Dictionary*. London: HarperCollins.
- Sternberg, R. J. 1987. Most vocabulary is learned from context. In *The Nature of Vocabulary Acquisition*, M. G. McKeown & M. E. Curtis (eds), 89–105. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Swanborn, M. S. L. & de Glopper, K. 1999. Incidental word learning while reading: A meta-analysis. *Review of Educational Research* 69: 261–285.
- Tomasello, M. 1992. *First Verbs: A Case Study of Early Grammatical Development*. Cambridge: Cambridge University Press.
- Tomasello, M. 2000. Do young children have adult syntactic competence? *Cognition* 74: 209–253.
- Tomasello, M. 2003. *Constructing a Language: A Usage-Based Theory of Child Language Acquisition*. Cambridge, MA: Harvard University Press.
- Wechsler, D. 1958. *The Measurement and Appraisal of Adult Intelligence*. Baltimore: Williams and Wilkins.
- West, R. F., Stanovich, K. E. & Mitchell, H. R. 1993. Reading in the real world and its correlates. *Reading Research Quarterly* 29: 34–50.

Appendix A: Definition task (Version A)

Choose the word that best matches the definition and write it in the blank. You can use the same verb more than once.

HOBBLE SAUNTER SCURRY SCRAMBLE STAGGER STRIDE SWAGGER BOLT
TRUDGE STRUT LIMP MARCH PACE PLOD PRANCE SIDLE SLINK AMBLE

1. _____: move hurriedly with small quick steps
2. _____: walk or more at a slow relaxed pace
3. _____: make one's way quickly or awkwardly up a steep gradient or over rough ground by using one's hands as well as one's feet
4. _____: walk in an awkward way, typically because of pain from injury
5. _____: walk with a stiff, erect, and apparently arrogant or conceited gait
6. _____: walk at a steady and consistent speed, especially without a particular destination and as an expression of one's anxiety or annoyance
7. _____: walk in a furtive, unobtrusive, or timid manner, especially sideways or obliquely
8. _____: walk slowly and with heavy steps, typically because of exhaustion or harsh conditions
9. _____: walk or move unsteadily, as if about to fall
10. _____: walk or behave in a very confident and typically arrogant or aggressive way
11. _____: walk with long, decisive steps in a specified direction
12. _____: walk in a military manner with a regular measured tread
13. _____: walk with difficulty, typically because of a damaged or stiff leg or foot
14. _____: move smoothly and quietly with gliding steps, in a stealthy or sensuous manner
15. _____: walk doggedly and slowly with heavy steps
16. _____: walk in a slow relaxed manner, without hurry or effort
17. _____: move with high springy steps; walk or move around with ostentatious, exaggerated movements
18. _____: run away suddenly out of control

Note: This version of the test contains definitions from the *New Oxford Dictionary of English*. The target responses are as follows: 1, scurry; 2, amble; 3, scramble; 4, hobble; 5, strut; 6, pace; 7, sidle; 8, trudge; 9, stagger; 10, swagger; 11, stride; 12, march; 13, limp; 14, slink; 15, plod; 16, saunter; 17, prance; 18, bolt.

Appendix B: Sentence completion task (Version A)

Below are 18 sets of sentences from which the verb has been removed. The sentences in each set originally contained one of the verbs from the list below. Can you guess what it is? Read all the sentences in the box first, then write your answer in the first blank, and continue to the next set.

Note: The sentences may require different forms of the verb (e.g. *amble*, *ambles*, *ambling*, *ambled*). You can use the same verb more than once.

AMBLE BOLT HOBBLE LIMP MARCH PACE PLOD PRANCE SAUNTER SCURRY
SIDLE SLINK SCRAMBLE STAGGER STRIDE STRUT SWAGGER TRUDGE

1. Missing verb = _____

The pig _____ into the undergrowth.

Pedestrians _____ for cover.

She _____ about the house picking up her children's toys where they had left them.

The mouse _____ across the floor and disappeared through a hole in the wall.

The noise of the explosion sent the villagers _____ back into their homes.

2. Missing verb = _____

The male bird _____ in front of the female.

The winner _____ forward to receive his prize.

This honour entitled her to _____ in front of the marching band at football games.

A peacock was _____ on the lawn.

The boys were _____ around trying to get the attention of a group of girls who were nearby.

3. Missing verb = _____

He _____ to his feet, swaying a little.

When he _____ in, they thought he was drunk till they saw the knife in his back.

We managed to _____ back up to the deck.

As we went into the bar, a drunken man _____ out the door.

Every morning she would wake up at 7 a.m. and _____ half-awake into the bathroom to get washed.

4. Missing verb = _____

I _____ up the stairs.

She _____ through blinding snow.

There was a stream of refugees _____ up the valley towards the border.

He _____ wearily along the path.

We _____ along the muddy track to the top of the hill.

5. Missing verb = _____

I _____ round the country roads for an hour.

He _____ into the foyer.

The pony _____ down the lane.

He _____ nonchalantly over to the phone.

She was just _____ along, going nowhere in particular.

AMBLE BOLT HOBBLE LIMP MARCH PACE PLOD PRANCE SAUNTER SCURRY
SIDLE SLINK SCRAMBLE STAGGER STRIDE STRUT SWAGGER TRUDGE

6. Missing verb = _____

He _____ off during Saturday's game.

The wounded soldier _____ along the road.

Two of the dogs were _____ badly.

Three minutes into the match, Jackson _____ off the pitch with a serious ankle injury.

Leaning on the old fashioned ebony cane she _____ across the floor.

7. Missing verb = _____

He was _____ around on crutches.

He _____ along as best he could.

The old man _____ past them.

Civilians and soldiers with missing legs _____ on crutches are a common sight.

The last time I saw Rachel she was _____ around with a stick, having injured her ankle skiing.

8. Missing verb = _____

There were a lot of people waiting to _____ aboard the small boat.

She _____ up the hillside and over the rocks.

We were _____ through the thick undergrowth when we suddenly came across a fast-flowing stream.

As the burning plane landed, the terrified passengers _____ for the door.

After waiting for over an hour, they _____ madly to get the best seats.

9. Missing verb = _____

Members of the Royal British Legion _____ past the Cenotaph.

They _____ through Norway.

Play a band and they begin to _____.

The soldiers _____ 90 miles in three days.

She _____ into my office demanding to know why I hadn't written my report.

10. Missing verb = _____

The pony was _____ around the paddock.

She _____ around the lounge impersonating her favourite pop stars.

When it was Vic's turn, he _____ about, lifting his knees high.

It's pathetic to see fifty-year-old pop stars _____ around on stage as if they were still teenagers.

I wish you children would settle down and stop _____ about.

11. Missing verb = _____

He _____ across the road.

He _____ confidently across the hall.

He _____ over the stream.

The soldiers _____ across the street with bazookas on their shoulders.

Clipboard in hand, she _____ purposefully up to the doors.

12. Missing verb = _____

We _____ up and down in exasperation.

She began to _____ round the office.

Alistair _____ up and down nervously, waiting for word from the surgeon.

By the time I arrived at the station, my father was already _____ up and down.

I hate to see animals _____ up and down in their cages.

AMBLE BOLT HOBBLE LIMP MARCH PACE PLOD PRANCE SAUNTER SCURRY
SIDLE SLINK SCRAMBLE STAGGER STRIDE STRUT SWAGGER TRUDGE

13. Missing verb = _____

The fox came _____ through the bracken.

All the staff have _____ off home.

I _____ away to my room, to brood in front of the fire.

The dog _____ out of the room with its tail between the legs.

He _____ away into the night.

14. Missing verb = _____

Look at that Charlie _____ down the street in his new suit!

The lord and his lady got up and _____ out.

They _____ into the room.

A group of young men _____ about outside the bar.

He _____ down the street after winning the fight.

15. Missing verb = _____

I _____ up to her.

She stammered some apology as she _____ towards the door.

A man _____ up to me and asked if I wanted a ticket for the match.

Tom _____ over to the pretty girl in the bar and asked if he could buy her a drink.

She _____ past him, pretending that she had not seen him.

16. Missing verb = _____

We _____ back up the hill.

The old man _____ along, hardly able to lift each foot.

We _____ wearily up the road carrying our heavy sacks.

We _____ through the mud.

Isn't it boring being a police officer, _____ along the streets all day?

17. Missing verb = _____

Adam _____ into the room.

All afternoon he _____ up and down, looking at the shops and the people.

He was whistling as he _____ along the beach.

He _____ by, looking very pleased with himself.

The children _____ down Sloane Street, loitering at the shop windows.

18. Missing verb = _____

She _____ for the door.

Passengers clearly overheard his shouted warning to the control room and they all _____ into the next carriage.

Frightened by the car horn, the horse _____.

He _____ blindly towards his father's fallen goat.

I was terrified that the horse would _____ and I would not know how to stop it.

Target responses: 1, scurry; 2, strut; 3, stagger; 4, trudge; 5, amble; 6, limp; 7, hobble; 8, scramble; 9, march; 10, prance; 11, stride; 12, pace; 13, slink; 14, swagger; 15, sidle; 16, plod; 17, saunter; 18, bolt.

Constructions and constructional meaning

Ronald W. Langacker

1. Introduction

Traditionally, a sharp distinction is drawn between an expression's syntactic structure and the lexical items it contains. The former constitutes its form, and the latter determines its meaning. This neat division of labor was continued in the generative tradition (at least in its early stages, e.g. Chomsky 1965), with its doctrine of autonomous syntax, the positing of discrete components, and its view of lexical items as clearly delimited elements. One issue that commonly arose in this perspective was whether a particular phenomenon was best handled "in the syntax" or "in the lexicon".

That question loses its force in constructional approaches, including both Construction Grammar and Cognitive Grammar, which hold that lexicon and grammar form a continuum of meaningful constructions (Fillmore 1988; Fillmore, Kay, and O'Connor 1988; Goldberg 1995; Michaelis and Lambrecht 1996; Croft 2001; Langacker 1987, 1991, 2000, 2005a). Nonetheless, a vestige of the question still lingers in the issue of whether certain aspects of clausal meaning are ascribable to the predicate or are solely due to the grammatical construction (Goldberg 1995; Langacker 2005b). The issue can only be resolved by considering a broad array of interrelated problems. My purpose here is to explore these in preliminary fashion and sketch a unified approach.

2. Categorization

Several major areas of cognitive linguistic investigation share the fundamental property of using networks as a basic mode of representation. The networks in question consist of conceptual structures linked by correspondences. In Construction Grammar, they take the form of intersecting hierarchies of lexicogrammatical constructions. Cognitive Grammar employs them for both category structure and the grammatical organization of expressions, characterized as assemblies of symbolic structures. Networks are further used for representing mental space configurations and conceptual integration (Fauconnier 1985, 1997; Fauconnier and Sweetser 1996; Fauconnier and Turner 1998, 2002). As special cases, conceptual integration (or blending) subsumes both metaphor and metonymy (Lakoff and Johnson 1980; Lakoff 1987; Turner 1987; Kövecses and Radden 1998; Panther and Radden 2004).

A unified approach to these varied phenomena is, I think, well justified. Of central relevance here is a common feature of these different kinds of networks that is usually left implicit despite its basic importance – namely, that the link between two structures in a network is often asymmetrical. One aspect of their connection is an evident *directionality*, whereby one structure has a certain cognitive priority vis-à-vis the other. In Construction Grammar, the networks are described as “inheritance hierarchies”, where one structure “inherits from” the other. In Cognitive Grammar, I speak of the “standard” and “target” of categorization, and analyze grammatical constructions in terms of “component” and “composite” structures. Metaphor is characterized as a mapping between a “source” domain and a “target” domain. Descriptive labels of the form “X for Y” reflect the directionality of metonymy. In conceptual integration, a “blend” is produced by projecting selected elements from the “input” spaces. Finally, a basic dimension of mental space descriptions is how the spaces are navigated, some providing “access” to others.

The directionality of a relationship has to be distinguished from the correspondences defining it. Sometimes there is no apparent direction. In blending, for example, the input spaces are connected to one another by correspondences, and both project to the blend, but there need be no asymmetry between them. Likewise, in grammatical composition the composite structure is accessed via the component structures, but neither of these is necessarily accessed through the other. And while lexical senses are often related by extension with a clear direction, there are also cases where they seem to be equal in status. For instance, if *hot* in the sense of ‘spicy’ represents an extension vis-à-vis the temperature sense, it is not evident that the latter’s application to either an experiencer (*I’m hot*) or an object (*The plate is hot*) has priority relative to the other.

Moreover, when a connection does exhibit directionality, the direction is not always irreversible. Although the primary directionality in blending is that of input structures projecting to the blend, it is also common for emergent features of a blend to be projected back to an input. Thus the blend in (1), projected from spaces representing the Titanic and American politics, is actually used as a comment on the latter. Another kind of example is back formation, e.g. the verb *belly-dance* deriving from the compound *belly-dancer*. In the original expression, *belly-dancer* represents the output, derived from *belly* and *dancer* in accordance with the noun compounding construction. But in the back-forming process, *belly-dancer* is one of the inputs, along with the suffix *-er* and the derivational construction for combining it with verbs. Running this in reverse, so to speak, yields the verb *belly-dance* as output.

- (1) *If Clinton were the Titanic, the iceberg would sink.*

In Figure 1(a), I represent the link between two structures, labeled X and Y. Dotted lines are correspondences, and a wedge indicates directionality. Diagram (b) shows an abbreviatory notation, where X and Y are merely labeled (rather than characterized), and correspondences are omitted. The arrow represents both the fact that X and Y are connected and the direction of their linkage.

Many instances of directed linkage are reasonably described psychologically in terms of X being used as a basis for apprehending Y. With respect to metaphor, it is often said that the source domain is used to understand the target domain. Categorization is likewise




Figure 1. Notations for directed linkage

a matter of apprehending the target in a particular way as specified by the standard. In such cases I will say that Y is *apprehended ás* X. This is a very general psychological phenomenon, something we engage in at every moment of our waking lives – we can't help it, and we can't function otherwise.¹ It happens when you see your mother and recognize her *ás* your mother. It happens when you hold a pen and recognize it *ás* a pen. It happens when you are hungry and experience it *ás* being hungry. It happens when you hear an English word and apprehend it *ás* that word (rather than just noise). In the broadest sense, it is simply a matter of activating established cognitive routines in dealing with current experience.

Crucially, it is not the case that apprehending Y *ás* X is merely the sum of X and Y. Hearing an utterance of *viagra* *ás* the word *viagra* is not reducible to the separate mental experiences of activating the stored acoustic image of that word plus hearing a stretch of raw, unanalyzed sound. At the very least, these experiences have to be *coordinated* as facets of a single, more complex experience in which the former is applied to the latter and used to interpret it. Generally this interpretive function serves to partially *constitute* the target, which can thus be recognized on the basis of only partial or degraded input. Probably you will recognize your mother if you only catch a glimpse of her under poor lighting conditions. Moreover, apprehension *ás* commonly produces *emergent* properties. This is evident in the case of metaphor, where using the source domain to understand the target domain results in metaphorical entailments. If we understand a theory *ás* a building, we can reason to the conclusion that too many ad hoc additions – tacked on in cantilevered fashion instead of rising directly from the foundations – might cause it to collapse.

From X and Y alone, therefore, the overall experience of apprehending Y *ás* X is not strictly predictable, owing to emergent properties as well as the possibility of being shaped by other inputs. Referring to X as the *standard* (S), and Y as the *target* (T), we can say that S and T contribute to the overall experience without being exhaustive of it. A full characterization must also include the details of their relationship and any other structures or properties that emerge.

We can distinguish several basic types of apprehension *ás*. One type is *full recognition*: the standard is fully manifested in the target with no significant distortion. For instance, you see a familiar face and it appears just as it always does. The standard and target are not equivalent: because S is a stored structure abstracted from previous experiences, T is

1. This is something so basic and fundamental that we need a convenient means of referring to it. The term I adopt here is *as*, pronounced with full stress and written with an accent: *ás*. This forms the basis for locutions like *apprehension ás*, *understand ás*, *recognize ás*, etc.




Figure 2. Kinds of recognition

usually more elaborate, i.e. specified in finer-grained detail. Still, matching the standard against the target reveals no discrepancy. Indeed, since *S* is fully and straightforwardly manifested in *T*, it is *immanent* in *T*, i.e. it “lies within” it. The cognitive processing which constitutes *S* is inherent in the more elaborate processing which constitutes *T*. As a consequence, full recognition engenders no awareness of anything other than *T*. The result of apprehending *T* *as* *S* – the overall experience residing in their coordinated activation – does not involve any content distinct from *T*. Since *S* fully meshes with *T*, and *T* is the structure being apprehended, *S* is effectively *transparent*.

This is hard to diagram, but I have made an attempt in Figure 2(a). As the structure being assessed, *T* is in the foreground, as indicated by the solid-line circle. As the basis for assessment, *S* is in the background, hence the dashed-line circle. *S* is wholly subsumed in *T*. Moreover, *S* is roughly coextensive with *T* in terms of the elements they invoke – the difference is not in their “coverage” but in their granularity (degree of specificity). I have tried to indicate this with the arrows from *S* to *T*, representing an *elaborative* relationship.²

We can speak of *partial recognition* in cases where there is some conflict between the standard and target; i.e. *T* matches *S* only when certain specifications of *S* are suspended. Apprehending *T* *as* *S* therefore involves some strain or tension. An example is recognizing a person you haven’t seen for many years – you recognize the face despite the effort or uncertainty caused by the age-altered features. With partial recognition, the conflict between *S* and *T* engenders an awareness of more than just *T*. Since the standard is not fully immanent in the target, its occurrence as part of the overall experience is not transparent: the discrepancy registered when *S* is compared to *T* constitutes an additional aspect of this experience.

With partial recognition, *S* and *T* can be related in various ways, sketched in Figure 2(b). The relationship can be *contrastive*: their content is basically the same, except

2. This could better be shown in a three-dimensional diagram, where *S* would be directly above *T* in a separate plane.

that they make inconsistent specifications with respect to certain features. For instance, your mother dyes her hair and you recognize her despite the difference in color. The relationship can also be *subtractive*, in the sense that an element of S is missing in T. For example, if you see a cat which lacks a tail you can still see it as a cat. The perception that something is missing resides precisely in the configuration shown: T is in the foreground, S in the background; when S is mapped onto T, they roughly coincide in terms of their coverage; however, some element of S has nothing to map onto.

An *augmentative* relation is more complex. The term is meant to indicate that some facet of S which T lacks is mapped onto it nonetheless. Instead of remaining in the background, as a virtual or missing element, it is projected onto the target to produce a new, augmented target T'. This is the hallmark of constitutive metaphor, where mapped features of the source domain are not perceived as being missing from the target, but rather as partially constituting the target. For example, we metaphorically construe a computer as a person by projecting onto it a mind, a will, and certain personality traits. It is not that we think of the computer as lacking these human properties – quite the contrary. It is by virtue of imbuing a computer with these traits that we describe it as being *stubborn*, *capricious*, or *vindictive*.

Finally, we can speak of *complex recognition* in cases where multiple standards are invoked for apprehending a single overall target. Each standard (S_1, S_2, \dots) is used to assess some facet of the target, as shown in Figure 2(c). The facets of T they assess (T_1, T_2, \dots) can be disjoint, overlapping, or coincident, and may or may not be exhaustive of T. An example of their being disjoint would be seeing a knife, fork, and spoon lying side by side. Collectively they constitute a single overall target – we apprehend them as a functional unit, the utensils that go with one place at the table. But at the same time we recognize them individually as a knife, a fork, and a spoon. An example where T_1 and T_2 overlap would be a single utensil with prongs at one end of the shaft and the bowl of a spoon at the other. We recognize both a fork and a spoon, but the two overlap in the target. An example of T_1 and T_2 coinciding is a spoon-like implement where instead of being rounded, the end of the bowl has prong-like projections. (Appropriately, this is sometimes called a *spork*.)

An instance of complex recognition is thus resolvable into a number of component relationships, in each of which a standard (S_i) is applied to some facet (T_i) of the global target (T). The global target can be thought of as an *augmentation* with respect to each local target, T_i . In principle, each component relationship can represent any of the types of recognition previously described.

In Figure 2, the spatial overlap of circles and ellipses represents shared conceptual content, while labels (S, T, etc.) indicate function. Directionality is implied by these functions but not explicitly shown. For analytical purposes it is helpful to adopt a network representation, as in Figure 3. Each node in the network is labeled for its function, and the arrows connecting them specify directionality. With this notation there is no explicit indication that the linked structures share conceptual content.

In augmentative recognition, features of the standard are projected onto the target, producing an augmented target (T') as the object of awareness. The arrow from S to T reflects S being applied to T in order to interpret it. The other two arrows indicate that both S and T contribute to T', and in that sense have conceptual priority. With respect to

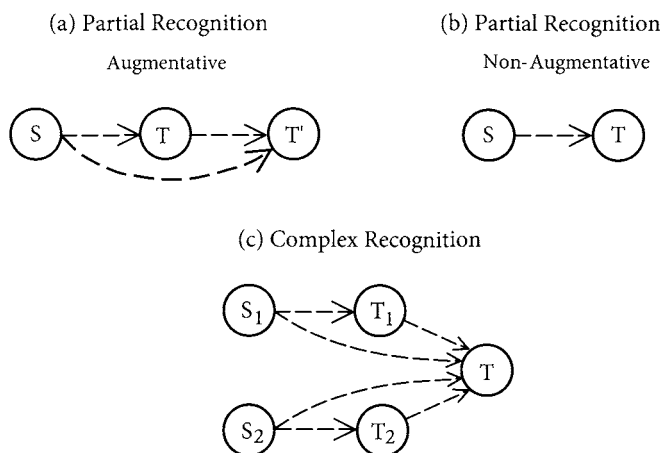


Figure 3. Network notation for types of recognition




Figure 4. Categorization compared to blending

augmentative recognition, the non-augmentative varieties (contrastive and subtractive) represent the special case obtained by collapsing T and T' . It is the case where S and T remain distinct in terms of their roles, i.e. no features are projected from S onto T . Hence a separate target, T' , fails to emerge. Complex recognition can also be viewed as a special case of augmentative recognition: the case where two judgments have the same augmented target, T . Since the component targets T_1 and T_2 are both facets of T , the latter amounts to an augmentation relative to each.

Augmentative recognition is itself an instance of blending. S and T are input spaces, and elements of each project to the blend, T' . This is shown in Figure 4(a), using the four-space model of Fauconnier and Turner (1998, 2002). The fourth structure, labeled G , is what they call the *generic* space. It represents the abstracted features shared by S and T . In Cognitive Grammar terms, it amounts to a *schema*, which both S and T elaborate. The solid arrows connecting G to S and T indicate these elaborative relationships.

The upper portion of Figure 4(a) is equivalent to diagram (b), used in Cognitive Grammar to represent an aspect of categorization, namely extension from a prototype. There has to be some basis for extension, something common to the standard and target which enables the latter to activate the former as the categorizing structure. This abstract-

ed commonality is a schema which the prototype and extension both instantiate. A dashed arrow connects the prototype and the schema to indicate that the schema emerges from the prototype, being induced by the process of extension. Categorization requires that the standard be recognized in the target. In the case of extension, where there is a conflict in specifications, a match is achieved only when certain specifications of S are suspended. It is the suspension of these discordant features that gives rise to the schema, which is fully recognizable in T. In sum, categorizing judgments involving extension amount to partial recognition. They may or may not result in a blend, T', obtained by projecting features of the prototype onto the target structure. Diagram (b) is neutral as to whether the partial recognition is augmentative (where T' emerges) or non-augmentative (either contrastive or subtractive).

If categorization by prototype amounts to partial recognition, categorization by schema amounts to full recognition, shown in Figure 4(c). The solid arrow indicates an elaborative relationship, where S is immanent in T. The configuration of diagram (c) can be seen as a special case of that in (a): it results when S is identical to G, i.e. no features of S need be suspended in order to arrive at a structure that fully matches T (S itself fully matches it). A consequence of S being fully subsumed in T is that a blended structure, T', cannot emerge. Any features that might be projected from S are already part of T, so T and T' cannot be distinct. In other words, full recognition is the special case of partial recognition where S collapses with G, with the consequence that T collapses with T'.

In this section I have sketched a unified account of various phenomena that are often considered separately. As previously (Langacker 1987: Ch. 10), I have attempted the integration of categorization by prototype and categorization by schema. I have related categorization to the general psychological phenomenon of recognition, or apprehension *ás*. Indeed, any attempt to distinguish them might prove arbitrary. This general phenomenon is in turn assimilated to mental space configurations. It represents the special case where one structure is applied to another in order to interpret it, thus giving rise to the standard/target asymmetry. When features of S are projected onto T, it also qualifies as blending (although blending does not necessarily involve S/T asymmetry). And among the instances of apprehension *ás* with blending are those we recognize as metaphor.

3. Composition

Grammatical composition is an integral part of this unified picture. A basic idea of Cognitive Grammar (henceforth CG) is that composition is a complex sort of categorization (Langacker 1987: 12.2). A grammatical construction is characterized in CG as an assembly of symbolic structures linked by correspondences and categorizing relationships. Because the symbolic structures qualify as mental spaces, it is also a mental space configuration. A construction can further be described in terms of blending, where the component structures function as input spaces, and the composite structure as the blend.

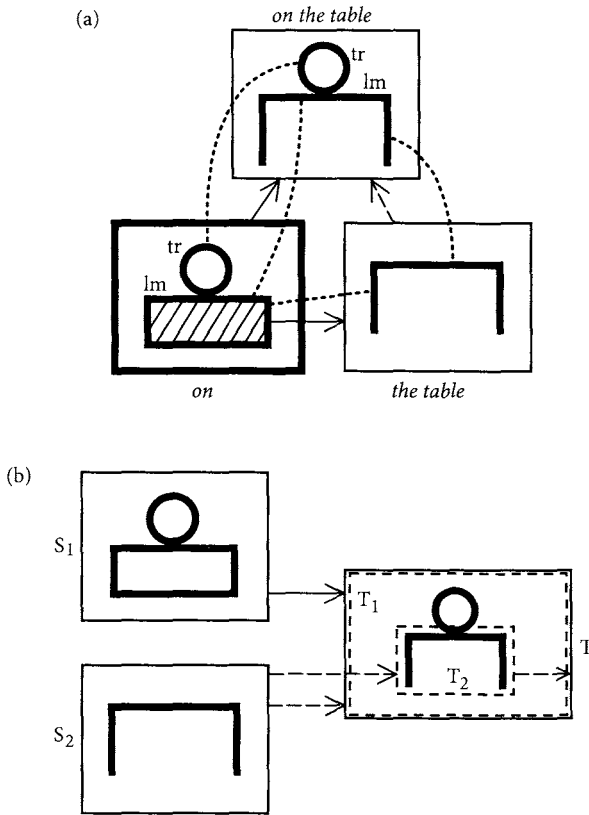


Figure 5. Constructions

Shown in Figure 5(a) is a typical CG representation of a canonical construction.³ The component structures are *on* and *the table*, and the composite structure is the prepositional phrase *on the table*. The component structures are connected (or integrated) through a correspondence equating the preposition's schematic landmark with the nominal profile. A solid arrow indicates that the nominal bears an elaborative relationship to the landmark (the preposition's elaboration site, marked by hatching). All the essential content of both components projects to the composite structure, which profiles the same relationship as the preposition. The preposition is thus the profile determinant (or head), as indicated by the heavy-line box enclosing it. The categorizing relationship between component and composite structures is one of elaboration in the case of *on*, extension in the case of *table*. Because they agree in profiling, *on* is schematic vis-à-vis *on the table*. With respect to *table*, on the other hand, *on the table* is an extension because they disagree in profiling.

Clearly, a construction like Figure 5(a) is an instance of complex recognition, as shown in diagram (b). The component structures function as the standards, S_1 and S_2 , and the

3. I will ignore elements like articles and tense, since our main concern is lexical content.




Figure 6. Composition

composite structure as the overall target, T . Dashed-line boxes indicate those facets of T which S_1 and S_2 serve to recognize, i.e. the local targets T_1 and T_2 . In this example, T_1 coincides with T : S_1 has full coverage in T , in that it schematically represents all the essential elements of the composite conception. By contrast, S_2 projects to just a portion of T (the relational landmark). So while S_2 contributes to the apprehension of T , with respect to T_2 the latter constitutes an augmentation.

A general scheme for composition is given in Figure 6(a). A and B represent the conceptual content of the two component structures. A' and B' represent the manifestations of A and B within T – their local targets of recognition. Relative to A and B , A' and B' are usually elaborated and sometimes distorted. C represents any further content, i.e. any portion of T not subsumed by A' or B' . In diagram (b) I give a simplified representation, which does not specifically indicate the local targets. The only arrows shown reflect the role of A and B in apprehending the overall target T .

The component structures should not be thought of as building blocks, but as stepping stones providing access to the composite conception. Rather than fully constituting the composite structure, they are better described as evoking it and imposing a particular way of apprehending it. For this reason I characterize their relationship in terms of categorization. But I also speak of categorization for the relation between the two component structures. In Figure 5(a), for example, the arrow indicates an elaborative relation between *on* (specifically, its landmark) and *the table*. It is unproblematic for a single structure (in this case *the table*) to serve simultaneously as the target in one categorizing relationship (vis-à-vis the other component) and the standard in another (vis-à-vis the composite structure). Notions like standard and target refer to functions rather than to structures per se. Hence the same structure can function in either capacity, or in both capacities when it participates simultaneously in multiple categorizing relationships.⁴

It may seem peculiar to say (with respect to their semantic poles) that *on* categorizes *the table* in the expression *on the table*. Formulated more precisely, however, the notion is fairly straightforward. It is merely being said that the preposition participates in the nominal's recognition (categorization in the broadest sense), and does so via a particular element, namely its schematic landmark (the elaboration site). In the phrase *on the table*,

4. This is comparable to the composite structure at one level of organization functioning as component structure at a higher level. Or for the same nominal referent to function simultaneously as the trajector of one profiled relationship and the landmark of another.

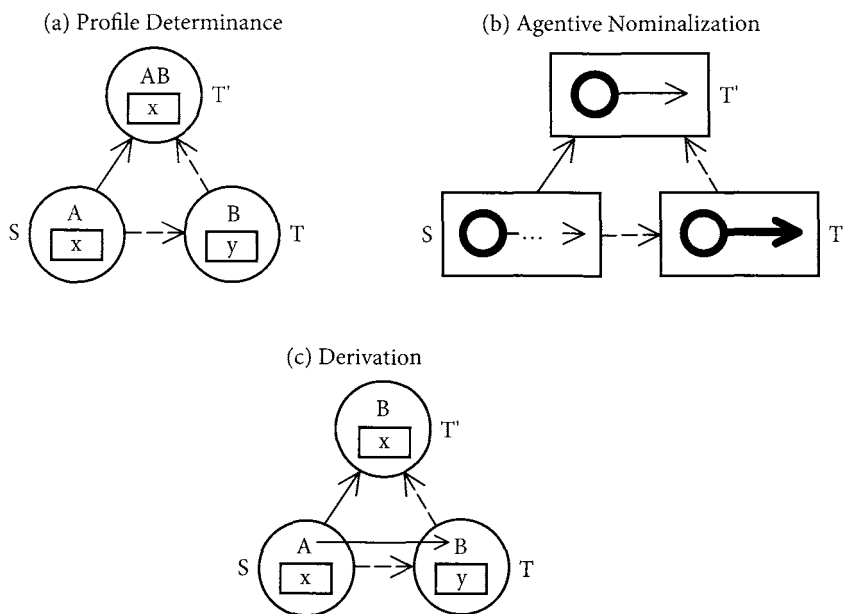


Figure 7. Profile determinance and derivation

the table is apprehended as the relational landmark, hence as a thing with certain properties (e.g. prototypically it has a surface).

But there is also a sense in which the preposition as a whole categorizes the object nominal. The preposition is the profile determinant, the component structure whose profile is inherited at the composite structure level. It can be described as projecting its relational profile onto the nominal, resulting in its content being apprehended as a relationship. The construction can thus be characterized as augmentative recognition, where *on* is the standard (S), *the table* is the target (T), and *on the table* is the augmented target (T'). Hence the construction has the configuration shown in Figure 3(a) and the lower portion of Figure 4(a). The composite structure T' blends the content of the components, S and T, which are asymmetrically connected in that S imposes its organization on T rather than conversely.

A construction of this sort is abstractly represented in Figure 7(a). For diagrammatic convenience, I am using capital letters (A, B, C, etc.) for conceptual content, and boxes labeled with lower case letters (x and y) for specifications pertaining to construal, notably profiling and trajector/landmark alignment. The diagram indicates that the augmented target T' inherits the content of both S and T, and construes it in the manner specified by S. Since the content and construal of S are both fully manifested in T', their relationship is elaborative (hence the solid arrow). Diagram (a) is a general representation of canonical constructions in which one component structure (S) functions as profile determinant. A specific example is *on the table*, shown in Figure 5(a).

A special case of profile determinance is derivation effecting a change of category. A stock example is the agentive suffix *-er*, which derives a noun from the verb stem it attach-

es to: *teacher, complainer, lecturer*, etc. As shown in diagram (b), *-er* evokes a schematic process as its base, within which it profiles the agentive participant. Elaborating this schematic process is the specific process profiled by the verb. And since *-er* is the profile determinant, the composite expression is a noun which designates the agent in that process.

Viewed in general terms, derivational constructions represent the configuration in Figure 7(c), which differs only slightly from 7(a). The primary difference between them is the solid arrow connecting A and B in diagram (c). It reflects the special property distinguishing derivation from other cases of profile determinance: the derivational element's schematic elaboration site – the substructure which corresponds to the profile of the other component – is exhaustive of its content. In terms of its semantic content, therefore, the derivational element as a whole is schematic vis-à-vis the other component structure. Its semantic contribution is limited to the profiling it imposes on the content supplied by the other component. The derivational element makes no independent contribution to the composite expression's content, precisely because A is immanent in B, hence totally subsumed by it. Thus the composite structure comprises the content of one component (B) construed with the profiling of the other (x).

Derivation would seem to be a clear instance of one component structure being used to apprehend the other. It is even clearer in examples of event nominalization, e.g. *occur* → *occurrence*, where the entire verbal process is reified and profiled as an abstract thing. A noun like *occurrence* is straightforwardly described as a process being apprehended as a thing. Constructions approximate this configuration to varying degrees. Derivation can thus be regarded as the limiting case where recognition as represents a construction's sole function. More broadly, as indicated in Figure 7(a), profile determinance can be characterized in terms of augmentative recognition being superimposed on the complex recognition known as composition.

Thus categorization (in the broad sense of recognition as) is equally characteristic of syntagmatic and paradigmatic relations. It is normally thought of as pertaining just to the latter. Whereas syntagmatic relations hold among the elements of complex expressions – a matter of how they combine syntactically with one another – paradigmatic relations hold among a category and its members, which function as alternatives to fill a given slot in a syntagmatic sequence. By contrast, CG views the syntagmatic and paradigmatic dimensions as overlapping and often non-distinguishable. They are seen as facets of the same global networks of directional relationships, where one structure motivates another or provides a means of apprehending it. A particular expression is motivated by relationships in both dimensions.

This blurring (or erasure) of the boundary between syntagmatic and paradigmatic relations follows from the notion of constructional meaning, especially as established by Goldberg (1995) in her seminal work. Goldberg showed quite clearly that constructions are independently meaningful. This is particularly evident in cases where an essential aspect of an expression's conceptual content is not inherited from the lexical items it contains. The verb *clear*, for example, does not intrinsically evoke the notion of a change of possession. Thus (2)a does not suggest that anyone else gains access to the desk by virtue of her clearing it. In (2)b, however, the recipient (*him*) does gain access to the entity cleared (*a place to work at her desk*). This is not due to the conventionalized meaning of

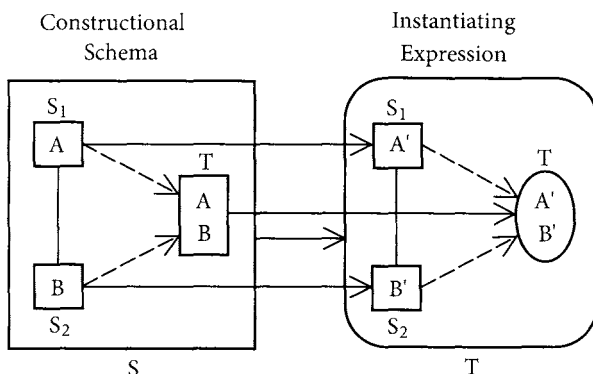


Figure 8. Instantiation of constructional schema

clear itself, but rather to its non-conventional use in the ditransitive construction. It is the construction, which includes it as part of its prototypical meaning, that contributes the notion of the recipient gaining access to the secondary object.

- (2) a. *She cleared her desk.*
 b. *She cleared him a place to work at her desk.*

In CG, grammatical patterns take the form of *constructional schemas*, i.e. schematized symbolic assemblies representing the abstracted commonality of instantiating expressions. Apart from their level of specificity, constructional schemas are precisely analogous to these expressions, consisting of component and composite structures linked by correspondences and categorizations. Those directional links are internal to a constructional schema, mirroring the ones internal to specific expressions which instantiate the schema. At the same time, the schema serves to categorize such expressions – to say that an expression represents a particular construction is to say that it is apprehended as an instance of it. We must therefore posit categorizing relationships at two levels of organization: those internal to a constructional schema or an instantiating expression, and those which hold between the schema and the expression.

Let us confine our attention to the simplest case, where the expression faithfully reflects the schema, and where the composite structure inherits all of its content from the component structures.⁵ The expression's categorization by the schema is then as shown in Figure 8. Unlike in previous diagrams, the difference between rectangles and closed curves (or boxes with rounded corners) is significant: the former indicate entrenched linguistic units, and the latter, structures which lack unit status. The constructional schema is shown on the left, the target expression on the right. Being an established unit, the schema is enclosed by a rectangle, as are the structures forming it. The overall expression, on the other hand, may well be novel, even if its components happen to be units (as indicated).

5. Usually the composite structure has content not inherited from either component. Instead of A'B', this structure would then be given as A'B'C, where C represents the additional content (cf. Figure 6).

Internally, the constructional schema specifies two component structures, with content A and B, both used to apprehend the composite conception, AB. The line connecting A and B stands for their *integration*: the correspondences and categorizations which specify the details of their combination to yield the composite structure. Within the expression, the component structures A' and B' are integrated in the manner specified by the schema. The composite structure A'B' is therefore related to A' and B' in the same way that AB is related to A and B. Solid arrows indicate the expression's categorization by the schema, i.e. its recognition as an instance of the construction. Globally, the expression (T) bears an elaborative relationship to the schema (S); hence the schema is immanent in the expression. This global relationship decomposes into several local ones: A, B, and AB are respectively elaborated by A', B', and A'B'.

The relationships internal to either the schema or the expression are syntagmatic, for they specify how simpler structures combine to form one of greater complexity. On the other hand, the external relationships would generally be regarded as paradigmatic: they hold between a class and the members of that class. The constructional schema as a whole defines a class of expressions. In local terms, A and B may themselves represent constructions, but they can also be lexical classes. A, for instance, might be schematic for the class of adjectives, and B for the class of nouns, the construction specifying a noun's modification by an adjective. A' would then be a specific lexical adjective, and B' a lexical noun. Even though both A' and B' are established conventional units, the combination A'B' could well be novel.

I am suggesting, however, that any sharp distinction between syntagmatic and paradigmatic relationships would be artificial. I have already pointed out that relationships in both dimensions are cases of apprehension as. I have further reiterated Goldberg's important observation that constructions are independently meaningful, and are often responsible for aspects of an expression's meaning which are not contributed by any component element. In examples like (2)b, a construction's semantic contribution is made visible by choosing a lexical item (*clear*) whose conventional meaning would not itself qualify it for use in the construction (ditransitive). There are also constructions whose semantic contribution is evident because they systematically specify meaning elements not inherited from either component structure. An example would be a possessive construction where the possessor and possessed are simply juxtaposed, with no morphological element (like English 's) to symbolize the possessive relationship.⁶ But even when the composite structure inherits all of its content from the components, the construction itself still makes an independent semantic contribution. In Figure 8, it is represented by the line connecting A and B: the constructional schema specifies how the content supplied by the component structures is integrated to form the composite structure. Since a different composite meaning emerges depending on how the components are integrated (e.g. through alternate correspondences or the imposition of alternate profiles), this information – generally not supplied by the components themselves – is crucial (Langacker 2003).

6. In such cases, the constructional schema itself specifies the additional meaning element at the composite structure level: if A and B are the component structures, the composite structure is ABC.

If we define component structures as those out of which a composite structure is assembled (in cases of full compositionality), it would therefore be arbitrary to exclude the constructional schema (even in terms of conceptual content). In Figure 8, the component structures contributing to the composite conception $A'B'$ include the constructional schema shown on the left, as well as A' and B' . All of these structures are evoked by way of apprehending the composite conception. Diagrammatically, this is reflected in the direction of the arrows. No matter where one starts, the paths defined by the arrows all converge on $A'B'$, the ultimate target within this multidimensional symbolic assembly.

4. Lexical description

If the same complex expression is used repeatedly, to the point of being entrenched in the minds of speakers and conventional in a speech community, it becomes a lexical item (defined in CG as any fixed expression). All the structures in Figure 8 then have the status of units, as shown in Figure 9(a). These conventional units include the composite conception $A'B'$, the target expression overall (T), and the expression's categorization by the constructional schema (S). An example is a word like *lecturer*, a well-established unit definitely recognized as an instance of the *-er* nominalization pattern.

Once an expression is established as a unit, it is susceptible to a gradual loss of *analyzability*. Intuitively, for example, I judge the expressions in (3) to be successively less analyzable. A novel expression like *antagonizer* is fully analyzable. At the opposite extreme, *ranger* comes close to being monomorphemic. The other lexical units are intermediate: I




Figure 9. Degrees of analyzability

am always aware of a *complainer* being someone who complains, but I do not always think of a *computer* as something that computes, or a *rover* as something that roves. The analyzability of frequently used expressions tends to diminish over time. Once the composite structure A'B' has the status of a unit, there is at least the potential for it to be activated independently of the structures and relationships – both syntagmatic and paradigmatic – that originally gave rise to it. Their lesser degree or likelihood of activation is indicated by the dashed-line boxes in Figure 9(b). Eventually, as shown in diagram (c), these motivating factors can fade away altogether, leaving A'B' as an isolated unit not decomposable into smaller symbolic elements.

(3) *antagonizer* > *complainer* > *lecturer* > *teacher* > *computer* > *rover* > *ranger*

Degree of analyzability is thus an important aspect of a lexical item's overall description. Along the paradigmatic axis, analyzability pertains to an expression's place in networks of constructions. Also figuring in the full description of a lexical item are categorizing relationships of another sort: those which hold among its alternate senses. These too can be modeled as a network of directed links, whereby one sense motivates another. For polysemy networks as well we can speak of analyzability and loss of analyzability.⁷

In cognitive linguistics, it is widely accepted that frequent lexical items are usually polysemous. Starting from an original or prototypical value, a lexeme tends to develop related senses by being extended to new kinds of situations. Originally applied to temperature, for example, the adjective *hot* has been extended to indicate spiciness. Multiple extensions, based on either the prototype or previous extensions, result in a network of related meanings that is often quite elaborate. Some extensions are clearly metaphorical, and others clearly metonymic. Of course, the two sorts of motivation are neither sharply distinct nor mutually exclusive. Do we refer to spicy food as being *hot* because the taste sensation is abstractly similar to the temperature sensation? Or is it because we feel hot when we eat it?⁸

When an extension first occurs, or is first learned by a given speaker, the expression's novel use is accompanied by activation of the basic sense that motivates it. The term *pig*, for example, is applied to a metal object sent through oil pipelines to test or clean them. When I first encountered the word *pig* used in this manner, I understood it as a metaphorical use of the animal name, and I continue to do so. But of course, if I worked in an oil field and used the term repeatedly every day, it would quickly lose its semantic analyzability for me – I would come to apprehend its extended meaning directly, without saliently or consistently activating its basic sense. Eventually its connection to the basic sense might be lost

7. The network metaphor may be overly discrete and should not be pushed too far (Allwood 2003; Zlatev 2003; Langacker 2006). All that really matters here is that there is often a direction of motivation in lexical meanings.

8. Queller (2001, 2003) argues against metonymy as a psychologically manifested mechanism of extension, suggesting instead that metonymic shifts result from discrepancies in an expression's contextual interpretation on the part of the speaker and the hearer. While this seems reasonable, we must also consider how later generations of speakers learn the polysemous lexeme. Exposed to both uses, it is plausible that a learner would apprehend one in relation to the other.




Figure 10. Loss of analyzability in metaphor

altogether, resulting in two words *pig* that are semantically unrelated, i.e. homophones. At the semantic pole, the overall development would thus be as shown in (4), where *PIG* represents the basic sense and *PIG'* the extended sense. Originally the latter is novel, hence only accessible via the former. Once *PIG'* and its relation to *PIG* become entrenched as units, the extended sense gains the potential to be accessed independently as the expression's meaning. Then, through a gradual fading of the motivating relationship – analogous to the development in Figure 9 – the semantic connection might be lost entirely.

(4) ([*PIG*] → [*PIG'*]) > [[*PIG*] → [*PIG'*]] > [*PIG'*]

As with composition, analyzability along this axis is a matter of degree. It is therefore not surprising that polysemy networks are often fuzzy. Even relationships that seem quite evident may not be accessed consistently or in all-or-nothing fashion. Nor is an extension necessarily motivated by just one established sense. So instead of being connected by specific and clearly discernible relationships, the senses in a network are generally better characterized as being motivated by other senses to various degrees (cf. Sandra and Rice 1995).

Of course, the fading of metaphors is not limited to extended lexical meanings. One reason the pervasiveness of conceptual metaphor is usually not appreciated is that the metaphorical apprehension of target domains becomes well-entrenched and independently accessible. For instance, we are so used to thinking of time *ás* money – i.e. as something we can *save*, *spend*, *borrow*, *lose*, *waste*, etc. – that we lose sight of the fact that such conceptions are metaphorically constituted. As shown in Figure 10, we can describe this as a loss of analyzability, where the blended structure *T'* is directly apprehended in its own terms, instead of emerging by actively invoking the source domain (*S*) to apprehend the target domain (*T*).

The two aspects of lexical description I have discussed so far – internal composition and polysemy – both involve directional relationships (categorization in a broad sense). The same holds for two further aspects, namely a lexical item's grammatical category and participation in grammatical constructions.

Rightly or wrongly, it is claimed in CG that basic grammatical categories have schematic conceptual characterizations, e.g. a noun profiles a *thing* and a verb profiles a *process*. Their conventional grammatical behavior, in a given language, is taken as being symptomatic of their meaning rather than definitional. Still, a lexical item's behavior is a crucial factor in ascertaining its grammatical category. Conceptual content does not come pre-identified as nominal content, verbal content, or whatever (though there are of course




Figure 11. Lexical representations and profiling

default expectations). Often the same lexical content supports the meaning of lexemes or alternate lexical senses representing different categories. In practice, then, a lexeme's category is ascertained by determining which constructions it participates in.

Various languages have been described as lacking grammatical categories at the lexical level. While this may be overstated (Haiman and Ourn 2003; Jacobsen 1979), let us accept its basic correctness as a starting point. In such languages, a lexical item's meaning presumably consists in just an array of conceptual content (cognitive domains or an ICM), with no particular facet of it being singled out for the kind of prominence – notably profiling – that determines grammatical category. Being neutral in this respect, a lexical item has the potential to be used in constructions characteristic of different categories. It is the very fact of its appearing in a construction that is responsible for a lexeme's (transient) apprehension as an instance of the category it specifies.⁹

For the sake of concreteness, consider the process of cooking. Our general conception of this activity, which I will refer to as the cooking scenario, is not inherently associated with any particular grammatical category. Of course, the scenario does make reference to a process, and to various things which participate in that process, including an agent and a patient. Depending on what is put in profile, therefore, it could supply the conceptual content for either a verb or a noun. Our assumption, though, is that a certain language invokes this scenario as the meaning of a lexical item without imposing any particular construal on it. Let us further assume, just for mnemonic convenience, that the form of this lexeme is *cook*. It can thus be represented as shown in Figure 11(a). As its lexical meaning, the form evokes a scenario whose essential content consists in a process with two central participants. It does not however impose a profile on this content, so it does not belong to any category defined in terms of profiling.

Suppose, now, that a speaker has occasion to use this lexeme as a noun. Specifically, it is understood as designating the agent when used in combination with a nominal grounding element, like an article or a demonstrative (Langacker 2004). In Figure 12(a), the box at the upper left stands for the relevant constructional schema, which specifies the combination of a grounding element with a noun to form a grounded nominal. To simplify the diagram, the composite structure is omitted; hence the inner boxes represent

9. This is quite comparable to "type shifting", discussed in Michaelis (2004).

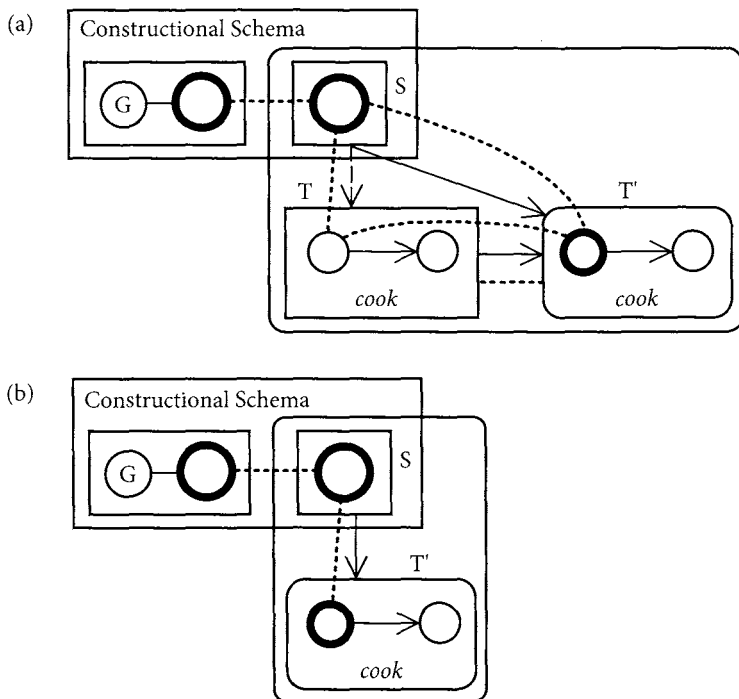


Figure 12. Categorization imposed by grammatical use

the component structures (the counterparts of A and B in Figure 8). Shown on the left, the grounding element situates a schematic thing (bold circle) vis-à-vis the *ground* (G), i.e. the speech event and its participants. The other component structure is the schematic characterization of a noun, whose profile corresponds to that of the grounding element (which profiles the thing it grounds).

The diagram shows the lexeme *cook* instantiating the noun slot in the constructional schema.¹⁰ Since *cook* is not intrinsically a noun, strictly speaking it does not satisfy the construction's specification in this regard; the dashed arrow indicates this discrepancy. Still, the nominal referent required by the construction is readily found. As shown by a correspondence line, it is identified with the agent in the cooking scenario. The relationship thus established between *cook* and the noun slot in the constructional schema amounts to *cook* being apprehended as a noun, and specifically as referring to the agent. The constructional schema imposes this transient categorization by projecting its profile onto the agent role in the scenario. The result of this projection is the augmented target shown at the right, which blends the scenario's content with noun schema's profile.

On the presumption that it is novel, this categorization is enclosed in a box with rounded corners. It can be described in various ways, which are non-contradictory in the unified perspective being offered. For one thing, it is a case of blending and augmentative

10. It is thus the counterpart of the relation connecting B and B' in Figure 8.

recognition, as indicated by the labels S, T, and T' (cf. Figure 3(a)). At the semantic pole, it is also equivalent to a case of composition with profile determinance (cf. Figure 7(a)). From this standpoint S and T are analogous to component structures, and T' the composite structure. The latter inherits its content from T and its profile from S, which is thus the profile determinant. Moreover, since S is schematic, contributing only its profile to the composite conception, the configuration is essentially a case of derivation (cf. Figure 7(b)). It is the grammatical construction, rather than a distinct morphological element, which fulfills this derivational function. This is a further instance of the blurring between syntagmatic and paradigmatic relationships.

In Figure 12(a), the standard, target, and augmented target are shown as separate boxes to indicate their distinct functions. But in terms of their content (and presumably in terms of the constitutive cognitive processing), they extensively overlap. In particular, T is wholly immanent in T', having the same content and being neutral with respect to profiling. They can thus be collapsed, as shown in diagram (b). The two diagrams are notational variants, the only difference being that (a) represents the lexeme more explicitly. On the other hand, (b) is more explicit in another respect: by showing more directly that the lexeme is apprehended as a noun in the context of this usage.

By assumption, the lexeme's apprehension as a noun is transient and context dependent. But it need not remain so indefinitely. Suppose speakers of this language have frequent occasion to describe a person as a *cook*. In this event, the entire configuration in Figure 12(b) would become entrenched and established as a conventional unit, as shown in Figure 13(a). One of the structures constituting this unit assembly is T', i.e. the lexeme *cook* apprehended as a noun. It is then just a matter of definition whether we want to say that the noun *cook* is a lexical item of the language. From the standpoint of CG, to exclude it from the lexicon would be arbitrary. After all, "the lexicon" – understood as a separate, discretely bounded component – does not exist. There are only symbolic assemblies, of various kinds, degrees of complexity, and levels of specificity. These assemblies are linked to one another by directional relationships, thus giving rise to more elaborate assemblies, with no clear distinction between the syntagmatic and paradigmatic axes. In this perspective, a lexical item can only be sensibly defined (if one chooses to retain the term) as a *fixed expression*, i.e. an actual expression (as opposed to a schema) established as a conventional unit. In Figure 13(a), *cook* is accorded this status, and since it profiles a thing, it qualifies as a lexical noun.

But it is a matter of definition. In the circumstance described, one might choose not to view *cook* as a lexical noun, on the grounds that its apprehension as a noun only occurs in the context of the nominal grounding construction. I believe this is ultimately misguided, since lexical categorization is never truly independent of occurrence in constructions. Be that as it may, we can certainly agree that in this circumstance *cook* is only a noun derivatively, by virtue of contextual augmentation.

Suppose, however, that this unit – *cook* apprehended as a noun – comes to be used independently of the nominal grounding construction. How might this happen? It could simply be a matter of the lexeme being used in other constructions characteristic of nouns (e.g. a plural construction, a modifying construction with an adjective, a noun compound construction). Each of these constructions would induce a nominal construal, in the same




Figure 13. Lexical categorization as loss of analyzability

way as the grounding construction. And each usage might be frequent enough for the overall configuration to become entrenched and conventionalized, as in Figure 13(a). Each of these would reinforce and further entrench the apprehension of *cook* as a noun. Moreover, from the standpoint of any particular construction, it amounts to a lesser degree of analyzability, as shown in Figure 13(b). That is, if the lexeme's apprehension as a noun is not limited to the nominal grounding construction, its occurrence is only occasionally accompanied by activation of that construction. The nominal sense of *cook* is neither consistently associated with the construction nor dependent on it. This loss of analyzability can eventually lead to the situation in diagram (c), where the nominal sense is accessible even in isolation.

At that point it would indeed seem arbitrary to deny *cook* the status of a lexical noun. The lexeme exists in two variants, one unprofiled and hence neutral with respect to category, the other a noun because it profiles a thing. We can further imagine that, in precisely analogous fashion, *cook* develops into a lexical verb by being used in constructions which impose a processual profile. The situation would then be as shown in Figure 11(b). In this hypothetical language, *cook* is polysemous, with nominal and verbal senses which each elaborate the more basic, neutral sense by imposing a particular profile on its content. From a usage based perspective (Barlow and Kemmer 2000; Langacker 2000), it is expected that many commonly occurring lexemes would develop one or more elaborated senses

representing particular grammatical categories. The language would thus have numerous lexical nouns and verbs, even granting that its basic strategy is for category-neutral lexemes to be freely used in either capacity.

How different, in this respect, would this hypothetical language be from English? Maybe not very. After all, English has both nominal and verbal variants of *cook*, and we can reasonably posit a schematized, neutral sense representing their abstracted commonality. The same holds for many other lexemes. English is quite accommodating, moreover, in permitting nouns to be used as verbs (Clark and Clark 1979), as well as other category shifts:

- (5) a. *Once again the delivery boy porched the newspaper.*
 b. *Procrastinate is something you never want to do.*
 c. *Having been both rich and poor, I can tell you that rich is better.*

The difference, then, may well be just a matter of degree – what proportion of lexical items have a primary categorization, and with what degree of freedom they can be used in constructions requiring a category they do not yet instantiate.

Even if lexical categorization is fairly rigid, it cannot be divorced from occurrence in constructions. Conceptual content alone does not make something a noun or a verb. It is only by construing that content in a certain manner – by imposing a certain kind of profile – that a lexeme is established as a category member. However, we cannot see, hear, or taste an expression's profile. In practice, we can only ascertain it by observing how a form is used, i.e. the constructions it occurs in. We learn that *cook* is a noun by observing its occurrence in constructions characteristic of nouns. This is so whether the categorization is primary (i.e. part of the lexeme's initially acquired value, as with *chef*) or represents a secondary development (as in our hypothetical example).

Since a lexeme is established as belonging to a grammatical category by virtue of occurring in constructions conventionally associated with it, occurrence in those constructions is one dimension of the lexeme's characterization. Let us once more consider Figure 13(a), representing the occurrence of *cook* in the nominal grounding construction. The square-cornered boxes indicate that the entire assembly (hence its component elements) has been established as a conventional unit representing a familiar usage. Because this assembly is the context required for *cook*'s realization as a noun, it is part of the overall description of the nominal variant. If there are multiple assemblies of this sort, all established as conventional units, each figures in the characterization of this variant. As shown in Figure 13, this can lead to the situation where *cook* is interpretable as a noun even in isolation. Its independent accessibility does not however imply that assemblies like the one in 13(a) disappear – they might very well remain as established units reflecting aspects of the variant's conventional usage. These assemblies, the *structural frames* a lexeme conventionally appears in, are part of its full characterization.¹¹ They show the arbitrariness of any strict division between lexicon and syntax, or between syntagmatic and paradigmatic relationships.

11. They are roughly comparable to the "strict subcategorization" features assigned to lexical items in the *Aspects* model (Chomsky 1965).

5. Some implications

I have outlined a unified approach to varied phenomena, ranging from grammatical composition to conceptual metaphor. It is based on directional relationships, manifesting processing asymmetries broadly describable as recognition, categorization, or apprehension *ás*.¹² Networks of directed relations permit a non-dichotomous account spanning such traditional distinctions as lexicon vs. grammar, derivation vs. extension, syntagmatic vs. paradigmatic relationships, and constructions vs. instantiating expressions.

Several points of clarification are needed concerning networks. For one thing, the network metaphor implies more discreteness than is sometimes warranted. It entails a specific inventory of nodes connected in a precisely determined way by clearly discernible categorizing relationships. I have already noted that this is unrealistic in the case of polysemy, and possibly in general; elsewhere (Langacker 2006) I suggest a less discrete alternative. Suffice it to say that the network metaphor must not be pushed beyond the limits of its usefulness.

One way in which the metaphor misleads is by inducing us to view the connected structures as distinct and non-overlapping, like beads on a string. They are not. In terms of their content, and no doubt in terms of neural circuits and cognitive processing, they extensively overlap, as I tried to indicate in Figure 2. It is only for analytical convenience (and to reflect their functional unity) that I pull apart these overlapping packages of content or processing activity and represent them in separate boxes and circles.

We must also bear in mind that the structures linked to form a network represent different levels of organization. In particular, a network configuration of any size – comprising structures and their connections – can itself be invoked in some capacity, coalesce as a unit, and function as a node in the network. As shown in Figures 9(a) and 13(a), complex structures of this sort emerge at multiple levels and participate in directed relationships.

Finally, a given structure participates in numerous relationships pertaining to various dimensions of linguistic organization. It can therefore function simultaneously in multiple roles, each as part of a different relationship. In morphological derivation, for example, the stem (e.g. *lecture*) is both the target with respect to the derivational element (*-er*), as shown in Figure 7(b), and a standard with respect to the composite expression (*lecturer*, as in Figure 6).

This unified perspective casts a different light on many traditional concerns. We have noted the absence of any definite distinction between syntagmatic and paradigmatic relationships. Another distinction that largely disappears is that of derived forms being lexically listed or constructed by rules. If these are seen as mutually exclusive options, either is problematic. Merely listing a form like *computer* in “the lexicon” fails to capture its relationship to the *-er* derivational pattern. On the other hand, deriving it by rule fails to accommodate the non-compositional aspects of its meaning (a *computer* is not simply ‘something that computes’). In the present perspective this problem never arises. Since component structures are used for apprehending the composite conception rather than

12. Though it may at some point be useful to draw distinctions, I am using these terms more or less interchangeably.

constructing it (functioning as stepping stones rather than building blocks), it is normal for a complex expression to incorporate conceptual content not inherited from either component, as shown in Figure 6. To be “constructed by rule” in this framework is to instantiate a constructional schema. When the instantiating expression is novel, as in Figure 8, the composite meaning can thus incorporate extracompositional features induced by the context.¹³ What it means for an expression to be “lexically listed” in this framework is that the entire symbolic assembly in Figure 8 – including both the expression and the sanctioning schema – achieves the status of a unit, as shown in Figure 9(a). Any extracompositional meaning consistently associated with the expression will be retained in the unit assembly. Moreover, the expression still instantiates the constructional schema and in that sense is still “constructed by rule”. It is just that activation of the entire assembly is an entrenched cognitive routine. Of course, the expression’s apprehension as an instance of the pattern tends to fade with time in the manner previously indicated.¹⁴

The question of lexical listing vs. construction by rule was originally tied up with the notion of *economy*, which in the early generative era was often reduced to the slogan “the shortest grammar is the best grammar”. From this standpoint, construction by rule was the preferred option, since listing resulted in longer descriptions and loss of generalization. The implicit assumption that only one option could be chosen is what I referred to as the “rule/list fallacy” (Langacker 1987: 1.2.3). In the present account, fixed expressions like *computer* are both listed as units and related to a general rule (constructional schema). The question might then be raised whether doing both ought to be resisted as uneconomical. After all, the configuration in Figure 9(a) appears to be quite redundant.

One response is that economy (in the sense of length of description, or “symbol counting”) has to cede priority to descriptive accuracy. If our concern is with how language is actually represented psychologically, then the appropriateness of an account combining listing and derivation is in principle an empirical matter. It would now be generally agreed, I think, that listing is necessary to handle the unpredictable idiosyncrasies of individual forms, and also that different forms have enough in common that generalizations need to be captured. The present framework does both by positing a constructional schema which individual expressions elaborate in alternate ways. Moreover, the apparent redundancy of Figure 9(a) is simply a matter of notation, where overlapping structures are represented separately for analytical convenience. What is actually being claimed (as indicated by the solid arrows) is that the schema is immanent in its various instantiations, as shown in Figure 2(a), and that these in turn thereby overlap, diverging only in the features that distinguish them.

For approaches that accommodate both lexical listing and participation in a general pattern, a related question then arises. One such approach is Construction Grammar, where constructions and instantiating expressions are connected in inheritance hierarchies. For

13. None is shown in Figure 8, which is limited to showing compositional aspects of the construction.

14. There are ways to handle the problem in other approaches. For instance, Jackendoff’s (1975) proposal of full lexical entries plus lexical redundancy rules is not drastically different from the present account. Arguably, however, the latter has the advantage that the problem never arises in the first place – an optimal description flows directly from basic principles rather than requiring special theoretical devices.

all intents and purposes, constructions and expressions are thus related as in Figure 9(a). The question that arises is whether “inherited” information comes to be associated with the expression only via inheritance, or whether the expression has it independently, sharing it with the schema. On the former account, the expression itself is only partially specified, consisting just of distinguishing information, while on the latter account it is fully specified, the information being represented (redundantly?) at both levels.¹⁵

To my mind, this is a pseudo-question. It stems from taking the network metaphor too seriously, in particular by viewing the connected structures as separate and distinct. The analytical convenience this affords comes at the cost of obscuring their actual relationship by showing the same content in different places in a spatialized representation. Though less convenient, a representation like Figure 2(a) might at least help avoid this erroneous entailment. As Goldberg correctly notes (1995:74), “a connectionist system can capture the redundancy without inefficiency by allowing inherited information to be shared information; that is, instead of stating the specifications twice, aspects of the patterns that are inherited are shared by two overlapping patterns”.

In view of the above, and for other reasons cited in the Construction Grammar and Cognitive Grammar literature (e.g. Fillmore, Kay, and O'Connor 1988; Langacker 2000), any definite boundary between lexicon and grammar has to be considered arbitrary. Rather than discrete components, they form a continuum of constructions (or symbolic assemblies). It should not be surprising, then, that it is sometimes hard to determine whether some aspect of an expression's meaning is due to the lexemes it contains or the construction they appear in. Lexical and constructional meaning overlap and are often non-distinguishable.

Certain aspects of this overlap are obvious, granted the meaningfulness of constructions and grammatical categories.¹⁶ In a clause like *Jill ate a juicy peach*, for example, the notion of transitivity is inherent in both the verb *eat* and the clausal construction. Likewise, the conceptual properties characteristic of adjectives are shared by *juicy* and the adjective slot in the schema for the noun-modifying construction. There are however aspects of the problem that are less obvious and more contentious. They pertain to how a lexeme comes to acquire its categorization, how many construction-related variants it has, and what happens when it is used in a construction where it does not fit.

In the previous section, I argued that conceptual content does not per se determine category membership. Given an array of content (such as the cooking scenario), a lexeme's grammatical category depends on how that content is construed, notably in terms of profiling and trajector/landmark alignment. What this means in practice is that category membership is indissociable from use in grammatical constructions. Being schematic in regard to content, the conceptual import of grammar consists primarily in the construal it imposes. In general, then, we initially ascertain a lexeme's category from the constructions

15. Analogously, Jackendoff (1975) considers both “full-entry” and “impoverished-entry” versions of his account based on lexical redundancy rules.

16. Cognitive Grammar accepts the meaningfulness of both. Construction Grammar is at best ambivalent about the latter (Langacker 2005b).

it regularly appears in. It is only through its grammatical behavior that the conventional way of construing its content is made visible.¹⁷

In terms of our previous example, the cooking scenario is not inherently either nominal or verbal, but holds the potential for either use. A lexeme that evokes it as its content is a noun or a verb depending on whether it is construed as profiling a thing (the most obvious candidate being the agent) or a process. Appearance in particular constructions, like the nominal grounding construction in Figure 13(a), is responsible for imposing a particular construal or confirming its conventional status. Thus *cook* is neither a noun nor a verb independently of its occurrence in nominal or verbal constructions. This is the basis for stating that unit assemblies like 13(a) – representing its conventional occurrence in structural frames (constructions) – is one dimension of a lexeme's overall description. To be sure, the number and inventory of these frames remains to be determined. It may well vary, and certain constructions are doubtless more important than others. Moreover, assemblies like 13(a) are entrenched to different degrees, and if enough of them help to establish a lexeme's categorization, it is possible that no single one becomes (or remains) a well-entrenched unit. A lexeme may then be activated as an instance of a category even without a structural context.

Nonetheless, occurrence in structural frames establishes a lexeme's categorization (reflecting the conventional construal of its content), and some array of frames figure (with varying degrees of salience) in its overall characterization. These frames are nothing other than the relevant constructional schemas. In Figure 13(a), the frame shown for *cook* is simply the constructional schema for nominal grounding. If it has the status of a unit, an overall symbolic assembly in which a lexeme appears in a construction is properly regarded as part of the characterization of each. I have emphasized the construction's role in effecting the lexeme's categorization. By the same token, in a usage based approach constructions are seen as complex categories, where constructional schemas coexist with subschemas spelling out their conventional range of uses. An assembly like 13(a), which incorporates a specific lexical item, occupies the lowest level in such a network. It is however the same assembly, whether we apprehend it as a facet of the lexical item or the global construction.

6. Lexical senses and constructions

Goldberg's important work (Goldberg 1995) deals extensively with the relation between lexemes and grammatical constructions. Her Construction Grammar account is broadly compatible with the one presented here, albeit with certain qualifications stemming from some differences between Construction Grammar and Cognitive Grammar (Langacker

17. This is oversimplified but is hopefully appropriate for the range of phenomena that most concern us. Other factors are of course involved. For example, we have default expectations (e.g. that a discrete physical object will be labeled by a noun) and are very good at intention reading (Tomasello 2003). We can therefore learn a noun in grammatical isolation if someone uses it while pointing to its referent. Appearance in constructions then serves to confirm this provisional learning.

2005a, 2005b). A similarity that ought to be noted is her attention to the question of how lexical items fuse with the constructions they appear in. It is the same question broached here in Figure 12, concerning a lexeme's apprehension as an instance of a category specified in a constructional schema. A main focus in what follows is a particular issue Goldberg raises in this regard, namely the number of construction-related senses a lexeme has.

Her basic stance, on grounds of parsimony, is that the number of senses should be minimized. Positing additional senses on the basis of occurrence in constructions is done with evident reluctance: "Although I have generally tried to avoid positing additional verb senses to account for each ... syntactic pattern, I do not rule out the possibility that *some* alternations must be accounted for by postulating distinct but related verb senses" (56). "[T]he semantics of ... the full expressions are different whenever a verb occurs in a different construction. But these differences need not be attributed to different verb senses; they are more parsimoniously attributed to the constructions themselves" (13). The arguments Goldberg advances do I think establish her point that many constructions construe a process in a way the verb itself does not. In such cases the discrepant aspects of meaning are ascribable only to the construction, not to the verb, with the consequence that a verb has fewer distinct senses than the number of constructions it appears in. What remains to be determined is the extent to which the polysemy of verbs should be minimized. What Goldberg apparently regards as the canonical situation is that a verb has just one construction-related semantic variant, so that all its other uses involve the imposition of constructional meaning.

The question I would raise is whether this situation should actually be considered canonical. The presumption that it is runs counter to usage based approaches and the prevalence of lexical polysemy, which is widely accepted in cognitive linguistics. In a way it represents the opposite extreme from the view that a verb always matches the semantics of a construction it appears in, so that the number of senses is equal to the number of constructions. Goldberg has successfully argued against the *maximal polysemy view*. I suggest that we must also resist the lure of the *minimal polysemy view*. Let me lay out the rationale for an intermediate position.¹⁸

We can start with Goldberg's example of *kick*, which occurs in all the constructions in (6). In the maximal polysemy view, its occurrence in these constructions implies that *kick* has eight different senses, distinguished by profiling, the choice of landmark, and salient evocation of entities specified by relational complements (such as *black and blue* and *into the stadium*). Hence the non-verbal elements in each sentence elaborate elements schematically invoked by *kick* itself, whose own semantics is always congruent with that of the clausal construction. By contrast, the minimal polysemy view holds that *kick* has only one meaning, presumably that of (6)a. In all the other expressions, departures from that single sense are imposed by the construction.

- (6) a. *Pat kicked the wall.*
 b. *Pat kicked Bob black and blue.*

18. I should emphasize that Goldberg does not represent the minimal polysemy view – at most she regards a single verbal sense as being canonical. What I am suggesting instead is that the canonical situation lies somewhere in the middle of the spectrum.

- c. *Pat kicked the football into the stadium.*
- d. *Pat kicked at the football.*
- e. *Pat kicked his foot against the chair.*
- f. *Pat kicked Bob the football.*
- g. *The horse kicks.*
- h. *Pat kicked his way out of the operating room.*

My own position is that the truth lies somewhere in between. I agree with Goldberg that we do not want to posit eight full-fledged senses for *kick* (and no doubt many more, when other constructions are considered). On the other hand, I find it dubious that only one has any cognitive status. For instance, its occurrence in the caused-motion construction, as in (6)c, seems quite familiar to me. Moreover, to say of a horse or a mule that it *kicks*, as in (6)g, is arguably a conventional pattern. For a basic verb like *kick*, it is also not implausible that an occasional occurrence in certain other patterns is enough to leave some cognitive trace. These matters can in principle be investigated empirically (e.g. in terms of corpus frequency). Here, though, I merely wish to clarify what is involved in properly formulating the question of how many senses there might be.

Agreeing that (6)a represents the basic sense of *kick*, I would say, then, that it also has a secondary caused-motion sense, i.e. that of (6)c, and that senses reflecting certain other uses may well be established to some degree (even if only slightly). We need to consider the arguments that might be advanced against this position and in favor of the claim – the minimal polysemy view – that *kick* has only the sense of (6)a. Three basic arguments advanced by Goldberg are parsimony, avoidance of circularity, and avoidance of implausible verb senses.

I have already suggested that the matter cannot be decided a priori on grounds of parsimony. While it is certainly more parsimonious to posit one sense rather than several, our objective is to model linguistic knowledge correctly, not to write the shortest grammar that generates all the grammatical forms.

What about the circularity (Goldberg 1995: 12) of “positing a new sense every time a new syntactic configuration is encountered and then using that sense to explain the existence of the syntactic configuration”? While that would indeed be circular, the move is only characteristic of the maximal polysemy view, not the intermediate position being advocated. I am not, for instance, claiming that (6)h establishes a special sense of *kick* that dovetails with the overall meaning of the *way* construction (Israel 1996). It is only when a certain usage is entrenched and conventionalized that a construction-specific meaning is said to emerge. (Even that will be qualified in what follows.) This is not in principle circular, even if in practice it may be hard to determine.

The same holds for the postulation of implausible verb senses. While the maximal polysemy view might force one to posit a *way*-construction sense of *kick*, or a ditransitive sense of *clear* for (2)b [*She cleared him a place to work at her desk*], the intermediate position imposes no such obligation. Suppose it is claimed that positing a caused-motion sense for *kick* would itself be implausible. I would argue that such a claim is gratuitous, since construction-related meanings (i.e. aspects of meaning pertaining to profiling, argument structure, etc.) cannot be ascertained independently of occurrence in constructions. The kicking scenario does not come with any pre-assigned category or grammatical

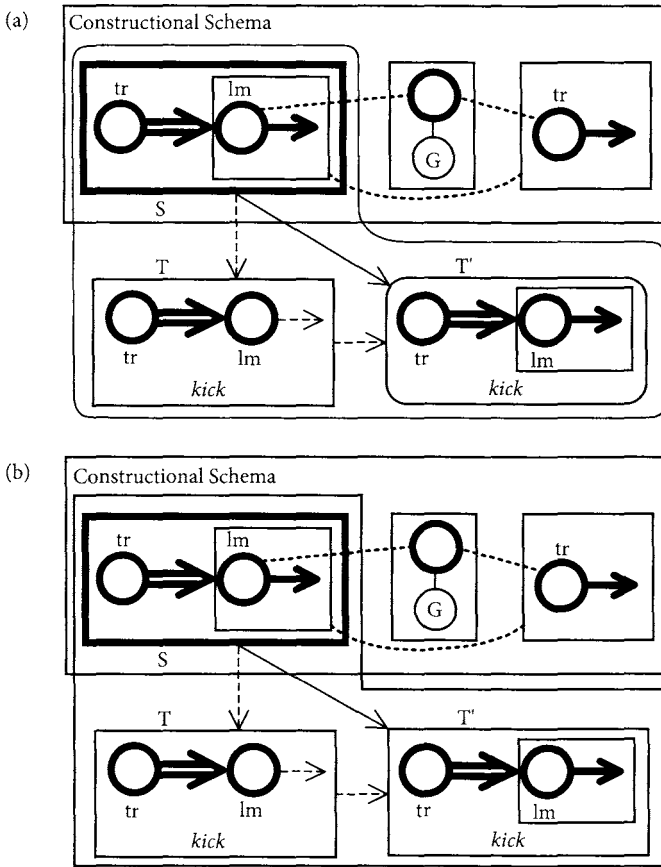


Figure 14. Induction of new lexical sense

frame. It lends itself to nominal construal (e.g. *kicker*) and to simple transitive use, but since kicking often results in motion on the part of the object kicked, it also lends itself naturally to the caused-motion construction.

Such arguments beg the question of what it means in the first place to say that a verb “has” a certain sense. In the perspective outlined here, it means that there exists an entrenched conventional unit where the form in question symbolizes a processual notion with the appropriate construal. Though part of the overall characterization, it does not matter if that construal is imposed by a construction or only occurs in the context of that construction. In the last analysis, all lexical senses reflect the constructions the forms appear in. It is the constructions that induce or confirm the construal which distinguish the kinds of senses being considered.

Suppose, for sake of discussion, that the simple transitive sense of *kick* is basic, so that initially its caused-motion use is incongruent with its meaning. The basic variant is represented at the lower left in Figure 14(a). It suffices to indicate that it profiles the transmission of force (double arrow) from trajector to landmark. A dashed arrow represents the potential of the landmark moving due to the force applied. While this is part of the kick-

ing scenario (one possible *dénouement*), it is not profiled by the verb and not necessarily evoked. The caused-motion schema, shown at the top, includes a verb, a nominal complement, and a relational complement describing a path of motion. The verb is schematic for predicates like *throw*, *toss*, *lift*, *pull*, etc., which profile an exertion of force that induces the landmark to move along some path. The nominal elaborates the verb's landmark, and the relational complement (typically a prepositional phrase) specifies the path.

Diagram (a) represents the use of *kick* in a sentence like (6)c, on the assumption that the simple transitive sense is the only one it has. It instantiates the verb slot in the caused-motion schema even though it does not fully satisfy the latter's specifications, for it does not profile, or even definitely invoke, the landmark's resultant motion. The categorizing relationship is therefore a matter of extension rather than elaboration. As an aspect of constructional meaning, the kicking is of course understood as causing the landmark to move. This fusing of the verb and the construction consists in *kick* being apprehended as a caused-motion predicate. Those specifications are projected onto the content of *kick* to derive the augmented target T', in which the kicking is specifically construed as inducing the landmark's motion.

In this particular example, the discrepancy between T and T' is fairly minimal. There is no change in trajector/landmark alignment, and the additional content put in profile is latent in the scenario supporting T. At least in terms of content, T is wholly subsumed in T'.¹⁹ Although the usage is novel, the extension is so natural that speakers will probably not even notice it. Originally, however, the caused-motion use of *kick* is based on and motivated by its simple transitive sense (by assumption, the only one it has). Should this usage recur, the entire configuration in Figure 14(a) will eventually be entrenched and conventionalized, as shown in diagram (b). At this point we can say that *kick* has a caused-motion sense: there is a conventional unit (T') in which *kick* specifically symbolizes the causation of motion by kicking. To be sure, this conventional understanding is based on the simple transitive sense and only occurs in the context of the caused-motion construction. But all extended meanings arise in the same general way, being understood in relation to the basic sense in a manner determined by the context. Thus if one accepts polysemy at all, it seems to me that *kick* has to be regarded as polysemous in 14(b). The relation between T and T' represents the first step in *kick* developing a polysemy network.

At this stage, the caused-motion sense of *kick* is still accessed via the basic sense. But once established as a unit, it has the potential to be accessed directly. As shown in Figure 15, we can thus anticipate a loss of analyzability, whereby *kick* comes to be used in this construction without the basic sense being evoked. The latter still exists, of course, and remains as the lexeme's prototypical meaning. The caused-motion sense is limited to the caused-motion construction, which in turn is part of its characterization.

Kick is usefully compared to *sneeze*, employed in Goldberg's now classic example:

(7) *He sneezed the napkin off the table.*

19. It makes little difference whether we think of T's profile conflicting with that of T', as shown, or being subsumed in it.




Figure 15. Loss of analyzability in lexical extension

I have described this case (Langacker 2005b) in terms precisely analogous to Figure 14. In its normal, intransitive use, *sneeze* does not meet the specifications of the caused-motion construction. When it is used in that construction, the force inherent in the process of sneezing is nonetheless construed as impinging on the landmark and causing it to move. Those aspects of its construal are imposed by the construction. If *sneeze* is substituted for *kick*, the fusion of verb and construction is just as shown in Figure 14(a).

Of course, the extension of *sneeze* to caused-motion construal is more drastic than in the case of *kick*. Being intransitive, *sneeze* has no focused landmark, and while it does involve the notion of force, the standard sneezing scenario does not include its being sufficient to cause the movement of objects. For this reason the example catches our attention – although it exploits something inherent in the basic sense of *sneeze* (T), the caused-motion sense (T') is strikingly different and unanticipated. Moreover, events of this sort are fortunately infrequent enough that the caused-motion sense is unlikely to coalesce as a unit (except among linguists who repeatedly cite this example). But if such events became commonplace owing to some cultural innovation – e.g. sneezing a napkin off the table

becoming a conventional means of expressing gratitude to one's host – *sneeze* could follow the path of *kick* in Figures 14(b) and 15.

We might also compare *kick* to *send*, which frequently occurs in both the caused-motion and the ditransitive constructions:

- (8) a. *He sent a birthday present to his daughter.*
 b. *He sent his daughter a birthday present.*

Send is clearly well established in both frames, implying two conventional senses differing primarily in their choice of landmark. It is not evident that either should be considered more basic than the other. Quite possibly, both variants of *send* are learned as such from the very outset. This is problematic for the minimal polysemy view, which requires that just a single sense be posited. The choice would seem to be arbitrary.

I suggest that *send*, *kick*, and *sneeze* represent three positions along a spectrum in regard to the status of construction-related variants. In the case of *send*, the caused-motion and ditransitive variants are presumably both well entrenched and more or less on a par. In the case of *kick*, the caused-motion sense is arguably conventional but no doubt secondary to the simple transitive sense. And for *sneeze*, only the intransitive sense is conventional – its caused-motion use is innovative. While they obviously need empirical confirmation, I consider these differences to be more or less self-evident. At least they are plausible examples of the kinds of situation one can reasonably expect to find.

Though presented in somewhat more detail, this is the same account I gave previously (Langacker 2005b) as an alternative to Goldberg's presumption of a single verbal sense being canonical. In their recent article on English resultatives, Goldberg and Jackendoff (2004: 534) allude to this account, but not altogether accurately. They ascribe to me the position that "the VP's complement structure ... is determined by the verb alone" in a range of constructions that include the ones in (9):

- (9) a. *I'll fix you a drink.* [ditransitive construction]
 b. *Fred watered the plants flat.* [resultative construction]
 c. *Bill belched his way out of the restaurant.* [way construction]

In other words, they impute to me the maximal polysemy view, which was not at all what was proposed. I noted in particular that (7) represents a case where the construction imposes its meaning, and that *sneeze* would have a meaning congruent with this construction only in the unlikely (though imaginable) event of this usage becoming entrenched and conventional. More fundamentally, I argued that lexical and constructional meaning are indissociable in those cases where a verb's occurrence in a construction is in fact conventionally established.

My basic objective was limited to suggesting that just a single verbal sense should not be accepted by default as the canonical situation. This was argued through specific examples, leaving open the question of the range of constructions which might induce new verbal senses, as in Figure 14. One option would be to make the strongest possible claim, consistent with the general picture outlined thus far. It might be proposed that, when a verb regularly occurs in a certain construction, it thereby acquires a sense congruent with the composite meaning of that construction. This would amount to a variant of

the maximal polysemy view, the difference being that the existence of a congruent sense is dependent on the verb being conventionally established in the construction.

Linguistic theorists generally have the attitude that “the strongest claim is the best claim”. However, one thing I’ve noticed about the strongest claims in linguistics is that they usually prove to be untenable. This one is no exception. To see why, let me introduce the notion of *skewing*.

Skewing is a discrepancy between a verb’s meaning and the composite meaning of an expression it appears in. A sentence like (7) represents a *skewing usage*: the verb *sneeze* appears in a construction despite a discrepancy between the process it designates (intransitive and non-causal) and the type of process profiled by the clause as a whole. The constructional schema specifies that the clause profiles the causation of motion, so in the context of this construction the event of sneezing is apprehended as an act of this sort. The term *skewing* is meant to capture the felt incongruity of this novel usage. But should the skewed usage become thoroughly familiar, it will no longer feel incongruous. As shown for *kick* in Figures 14 and 15, the caused-motion construal originally imposed by the construction comes to be established as a conventional value of the lexeme itself.

This putative development prompts consideration of the claim that a verb which regularly occurs in a certain construction thereby acquires a sense congruent with the composite meaning of that construction. But in fact this claim is too strong. To see this, we need only consider a simple passive, e.g. *Pat was kicked by a mule*. This is also a case of *skewing*, even though the verb and the clause profile the same event. They are discrepant in regard to trajector/landmark alignment, i.e. the choice of primary and secondary focal participants: the verb confers trajector status on the agent (with the patient as landmark), while the clause confers it on the patient (thereby defocusing the agent). Here, though, we do not want to say that *kick* itself develops a passive meaning just by virtue of occurring in the passive construction, however frequent this might be. Even if the passive locution *be kicked* becomes thoroughly entrenched and conventional, it would seem completely implausible to say (in accordance with the claim) that the verb itself has a distinct sense in which the patient is focused as trajector.

Why, then, are the passive and caused-motion constructions different in this respect? The difference is that the passive construction specifically incorporates a *skewing element*, which I identify as the passive participial inflection. The specific function of *-ed* is to reconstrue the process designated by the verb stem (in this case *kick*) by conferring trajector status on the patient rather than the agent.²⁰ In other words, *skewing* is inherent in the passive construction itself, as opposed to being induced by a mismatch between a verb and the constructional position it occupies. The constructional schemas for both the caused-motion and the passive constructions incorporate a verb as one component structure. The *skewing* in (7) is due to *sneeze* conflicting with the specifications of this component. By contrast, *kick* is in full conformity with the verb stem of the passive constructional schema. And since *kick* “fits” the construction perfectly, its passive use provides no basis for semantic extension.

20. It further construes it in summary rather than sequential fashion; *be* reimposes sequential scanning to derive the process profiled by the clause as a whole (Langacker 1982).

In addition to skewing usages and skewing elements (like *-ed*), we can recognize *skewing constructions*. These represent the special case where the skewing element is the construction as a whole, rather than any particular component. An example is the *way* construction. As described by Goldberg (1995: Ch. 9), there are two basic variants: one where the verb specifies the means by which movement along a path is brought about, e.g. *kick* in (6)h, and one (less typical) where the verb specifies the manner of moving, including activities that accompany the movement without being intrinsically related to it, such as *belch* in (9)c. While the composite expression profiles the process of moving along a path, neither *way* nor the path-specifying complement is processual. The verb does profile a process, but in general it is not one involving translational motion. The profiling of movement along a path is something that emerges at the composite structure level, a function of the construction as a whole rather than any specific element. It is thus a skewing construction because, by its very nature, the construction induces a discrepancy between the meanings of the verb and of the composite expression. When used in this construction, a form like *kick* or *belch* therefore conforms to its expectations in regard to the verb. Hence even conventional use in the *way* construction provides no impetus for the verb to develop an extended sense.

If the strong claim is erroneous, what would constitute a weak but possibly tenable claim? It would simply be that an element which regularly fills the verb position in a construction tends to develop a meaning congruent *with that position*. But this is almost a truism. It amounts to a special case of the platitude that an expression's meaning is shaped by how it is used. Of course, it is hard to find persuasive evidence that a new lexical sense has or has not developed when that sense is induced by and thus limited to a particular construction; since it does not result in a change in form, the process in Figures 14 and 15 is basically invisible. Preliminary investigation has convinced me that finding clear arguments bearing on the matter is very difficult owing to the multiplicity of interacting factors involved.²¹ More basically, however, the issue has no real substance without some prior understanding of what it is for a lexical item to have a certain sense. I have tried to give a principled explication of this notion, one situated in a broader perspective and an integral part of the unified account it affords.

7. Derivation by construction

Skewing constructions are further indication of the blurring between syntagmatic and paradigmatic relationships. They are cases where a constructional schema – an abstracted categorizing structure – performs a derivational function comparable to that effected by a specific morphological element, like the passive *-ed* or a causative suffix. Certain aspects of the content or construal of instantiating expressions are imposed by the construction itself, as opposed to being inherited from any component structure. In effect,

21. These include the range of potentially relevant constructions, the specifics of their analysis, the possibility of reanalysis, problems in assessing what putative “tests” actually show, and the many considerations that influence well-formedness judgments.

the constructional schema serves as an additional component structure, usually with the role of profile determinant.

From the CG perspective, constructions are able to fulfill this function because component structures serve as stepping stones for reaching the composite structure, rather than as building blocks for constructing it. There is no requirement or expectation that a composite structure inherit its construal or all of its content from the component elements. A construction is just an assembly of symbolic structures linked by correspondences and categorizing relationships. Since the component structures merely categorize and motivate the composite structure (rather than constituting it), there are usually aspects of the composite meaning which are not contributed by any component. In some cases the discrepancies themselves follow a regular pattern. They can then be ascribed to constructional meaning as represented in the constructional schema. Despite the discrepancies, instantiating expressions can therefore still be fully compositional given the standard definition: an expression is said to be *compositional* when the meaning of the whole is regularly determined by the meanings of its parts in accordance with a *rule of semantic composition*. In CG, a rule of semantic composition is simply the semantic pole of a constructional schema.

This flexible view allows a straightforward treatment of diverse constructions. Let me briefly survey some representative examples. In one class of cases, a clause – which in CG is claimed always to profile a process – does not contain any verb to supply one. Naturally, the understood process is quite minimal in terms of content. Typically it is merely an identity or reference point relationship (Langacker 1999). Sentence (10) illustrates an identity construction in Luiseno (a Uto-Aztecan language of southern California). Overtly it consists of just two nominals, which are simply juxtaposed. Semantically it specifies their referential identity.

- (10) *Xwaan no-kaytu.*
 Juan my-enemy
 'Juan is my enemy.'

Where does this identity relationship come from? One option is to posit a zero verb. But while I do not claim that such a move is never justified, more is needed by way of justification than the tacit assumption that all facets of the composite meaning have to be inherited from component structures. An alternative, sketched in Figure 16, is that the identity relationship (represented by a double bar) emerges at the composite structure level as an aspect of constructional meaning. More precisely, it is profiled by the composite structure of the constructional schema, which specifies just two component structures, both nominals. In a sentence like (10), the two nominals elaborate these schematic components, and since they “fit” the slots they instantiate, there is no inducement for them to develop a relational meaning congruent with the overall construction.

In another class of cases there is a verb, but its trajector does not correspond to the clausal subject. The expressions in (11) illustrate one well-known construction of this sort:




Figure 16. Identity construction

- (11) a. *The garden is swarming with bees.*
 b. *My cat is crawling with fleas.*
 c. *The whole town rang with church bells.*
 d. *The sky exploded with fireworks.*

It is actually the bees that swarm, the fleas that crawl, the bells that ring, and the fireworks that explode. The effect of this activity is to create the perceptual impression of an entire location – the garden, cat, town, or sky – being suffused with it, and thereby being the locus of a visual or auditory experience (cf. Dowty 2000). The verb appears in its basic form, i.e. there is no derivational element (comparable to a passive marker) which imposes a different choice of trajector on its processual content. This skewing is effected by the construction as a whole.

The relevant constructional schema is sketched in Figure 17. The two component structures are the verb, shown at the left, and a *with* phrase, shown on the right. The dashed-line box around the verbal process represents the location where the process occurs. While a location is implicit in the meaning of *swarm*, *crawl*, etc., the verb itself leaves it in the background (hence the dashed lines). On the other hand, the *with* phrase makes it prominent as trajector of the profiled relationship. The sense of *with* invoked in this construction is the one which figures in expressions like *the desk with a vase on it*, *the man with a wig*, or *a girl with freckles*: it profiles the relationship between a location and something found in that location.

Two correspondences effect the integration of the verb and the prepositional phrase. First, the verb's trajector is identified with the phrase's landmark (specified by the prepositional object). Second, the trajector of the prepositional phrase – the location where the landmark is found – is identified with the implicit location hosting the verbal process. As a consequence, that location hosts both the referent of the prepositional object and the activity it engages in. This provides the composite structure's essential content. In terms of construal, it inherits its processual nature from the verb and its choice of trajector from the prepositional phrase. The composite expression (e.g. *swarm with bees*) thus designates the process of its trajector (a schematic location) hosting an activity (*swarm*) on the part of

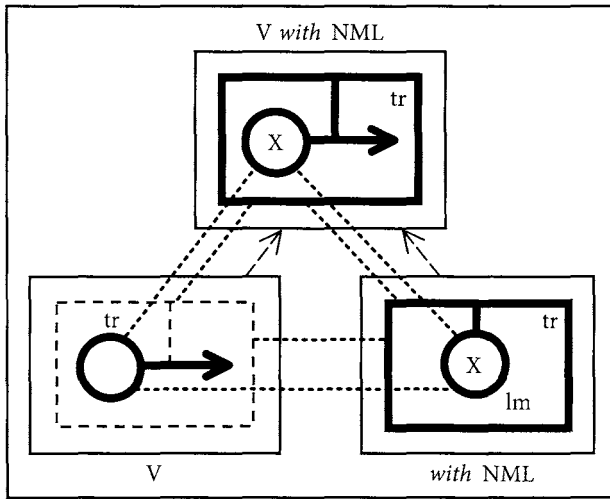


Figure 17. Location-subject construction

the prepositional object (*bees*). At a higher level of organization, the processual trajector is specified by a subject nominal (*the garden*) construed as designating a location.²²

This is a skewing construction, since the composite structure does not fully preserve the profiling and trajector/landmark alignment of either component. A verb like *swarm*, *crawl*, etc. does then conform to the specifications of the schematic verb it instantiates. Thus use in this construction is not per se expected to induce a new verb sense. Possible confirming evidence comes from the contrast between (12)a–b on the one hand, and (12)c–d on the other. Let us suppose (or imagine) that *swarm* and *crawl* are well-enough entrenched in this construction that a new sense could emerge if one were going to. If there were such senses, we might then expect them to support the derivation of the adjectival forms *swarming* and *crawling*, analogous to *shouting* and *screaming*, but these do not seem plausible, even in a context where the relevant senses ought to be salient. I realize, however, that this is anything but an air-tight argument.²³

22. Certain aspects of the composite structure's content and construal are emergent, rather than being inherited from either component structure. These include the notion of the activity suffusing the location, and the salience of the perceptual experience it engenders. I have not attempted to represent these in the diagram.

23. For example, one could counter that *swarm* and *crawl* do have the relevant senses, but only when specifically used in the construction; hence they are not available for adjectival derivation. One verb that does appear in adjectival form is *teem*: *The woods are teeming with wildlife. The teeming woods ...* This verb is not however very common in American English, and some American English speakers are not too sure what it means. (My dictionary defines it as follows: "be full of or swarming with".) Even so, it may indicate that there is no absolute prohibition against a new verb sense being induced to make it congruent with the overall expression, even with a skewing construction.




Figure 18. A resultative construction

- (12) a. *Her child was shouting with joy. The shouting child ...*
 b. *Her child was screaming with pain. The screaming child ...*
 c. *Her garden was really swarming with bees. *The swarming garden ...*
 d. *Her cat was crawling with fleas. *The crawling cat ...*

Other cases of skewing leave the trajector unchanged but affect the nature of the profiled process and the inventory of non-subject complements. Examples of this sort include the various kinds of resultative constructions discussed by Goldberg and Jackendoff (2004). We will only consider the type illustrated by (9)b, *Fred watered the plants flat*, which they describe as follows:

- (13) Syntax: $NP_1 V NP_2 AP_3$
 Semantics: $X_1 CAUSE [Y_2 BECOME Z_3]$
 MEANS: [VERBAL SUBEVENT]

In (9)b, Fred causes the plants to become flat by means of watering them. This is a skewing construction because normally the verb *water* is not construed as causing a change of state (at least not the one invoked by the adjective *flat*).

The constructional schema for this type is shown in Figure 18. The notation for the verb merely indicates that the trajector interacts with the landmark in some manner. In the notation for the adjective, the box represents a property (p), and the line connecting the box to the circle indicates that the trajector exhibits this property. Observe, now, that at the composite structure level this line is replaced by an arrow. The import of the arrow is that the participant corresponding to the adjective's trajector, rather than simply exhibiting the property, undergoes a change of state which results in that situation. This difference is one discrepancy between the component and composite structure levels. The other discrepancy, represented by the double arrow, is the notion of causation: the trajector-landmark interaction is construed as causing the landmark's change of state. Thus, in

addition to the interaction profiled by the verb, the relationship profiled by the clause as a whole includes both causation and the change of state that results.

The disparity that Goldberg and Jackendoff refer to as “syntax” vs. “semantics” does not support the autonomy of syntax or the need for distinct syntactic and semantic components. As shown in Figure 18, the construction is readily described in CG, which posits only symbolic assemblies for the description of lexical and grammatical structure. Semantics is simply one pole of such assemblies (the other being phonology). The formulas in (13) correspond to different facets of the symbolic assembly in Figure 18. The formula labeled “semantics” reflects the composite structure’s semantic pole. The formula labeled “syntax” represents the semantics of the component structures (since grammatical category is determined by an aspect of their meaning, namely profiling) as well as their phonological arrangement at the composite structure level. Basically, then, the disparity merely indicates the derivational function of a skewing construction.

As a final case, I want to consider an interesting twist on the notion of skewing constructions, one that highlights the derivational function and the absence of any clear boundary between syntagmatic and paradigmatic relationships. It is illustrated by the verb from our previous example, namely *water*, in a basic, non-resultative use, as in (14)a. But is *water* really a verb in the first place? It is also a noun, and if forced to choose one would certainly judge its nominal sense to be more fundamental. Conceivably one might argue that *water* is lexically listed only as a noun, and that its construal as a verb is an aspect of constructional meaning. It would thus be a skewing *usage*. Enhancing the plausibility of this claim is the observation by Clark and Clark (1979) that virtually any noun can be put to use as a verb, even a seemingly unlikely candidate like *porch* or a proper name like *Houdini*. On this account the shift to verbal use is effected without any separate derivational element – it is a function of the usage itself.

- (14) a. *Fred watered the plants.*
 b. *The delivery boy porched the newspaper.*
 c. *She had to Houdini her way out of the closet.*

I believe this is what we do in fact want to say for examples like *porch* and *Houdini*. They are simply learned as nouns, and English is permissive in regard to letting nouns be used as verbs. But this hardly seems plausible in the case of *water*, whose verbal use is thoroughly familiar and perfectly conventional. When a noun’s apprehension as a verb becomes entrenched and conventionalized, it thereby has a verbal sense. The verbal sense of *water* is so well established that its occurrence in (14)a no longer represents a skewing usage.

Clark and Clark point out that novel uses of this sort are extremely varied and dependent on presumed mutual knowledge for their specific interpretation. For instance, we depend on a familiar newspaper delivery scenario when we interpret (14)b as meaning that the delivery boy managed to throw the paper onto the porch (and not, say, that he fashioned it into a model of a porch). There are however some basic patterns in which the referent of the verbalized noun regularly has a certain kind of role vis-à-vis the clausal object. *Water* instantiates a pattern where the source noun is something *applied to* the object. Some other instances are listed in (15)a. In another pattern, source nouns like those

in (15)b designate something *removed from* the object. In the pattern of (15)c, the source noun indicates the resultant *location* of the object, as in the *porch* example. And so on.

- (15) a. *water, salt, pepper, butter, saddle, tile, seed, fuel, cork, asphalt, soap, wax, thread ...*
 b. *peel, pit, core, milk, juice, weed, skin, fleece, scalp, gut, scale, shell, husk, dust ...*
 c. *pocket, jail, warehouse, bench, corral, stable, file, package, bag, bottle, can ...*

While the noun use in each case seems more basic, the forms in (15) are all well established as verbs. Arguing for their lexical status is the fact that they are conventionally interpreted in a specific way which is less than fully predictable if not idiosyncratic (e.g. in American English usage, you normally *can* tomatoes by putting them in a jar). The emergence of new lexical senses is precisely what we expect with a skewing usage. If *porch* is not yet a lexical verb, it has the potential to become one through the continued occurrence of expressions like (14)b.

In English, the extension of nouns to verbal use is not limited to a single construction; e.g. in (14)c it happens in the *way* construction (where *elbow* and *claw* are lexicalized examples). The denominal verbs in (15) all occur in simple transitive clauses. Their emergence was no doubt facilitated by the existence of numerous underived transitive verbs corresponding to each pattern. The verbs in (16)a–c correspond to the respective patterns in (15): in (a) something is *applied to* the object; in (b) something is *removed from* the object; and in (c) the object ends up in a new *location*.²⁴

- (16) a. *feed, paint, fill, stuff, cover, wrap, dress, clothe, decorate, baste, fertilize ...*
 b. *clean, wash, scour, clear, empty, drain, shave, rob, mine, deplete, eviscerate ...*
 c. *move, throw, raise, lower, drop, sink, import, remove, plant, bury, exhume ...*

We can reasonably suppose, therefore, that each recurrent pattern of nouns being used as verbs invokes a low-level constructional schema where the verb is schematic for verbs of a certain sort: verbs of application, removal, relocation, etc. Making this assumption, and further assuming that the usage is novel, the verbal use of *porch* is sketched in Figure 19. The constructional schema is a special case of the schema for simple transitive clauses. It is more specific than the general schema because the verb profiles the action of moving the landmark to a new location (represented as a rectangle). As an instance of skewed usage, the noun *porch* instantiates the verb slot in the construction. The categorization is of course one of extension rather than elaboration, since the nominal profile of *porch* is incongruent with the schematic verb's processual profile. It is nonetheless effected by virtue of partial conceptual overlap: the nominal profile corresponds to the verbal landmark's new location. As a consequence, *porch* is apprehended as a verb of relocation in the context of this construction. There is no need to posit a zero derivational morpheme – the usage itself serves a derivational function.

The example is quite analogous to the caused-motion use of *kick*, diagrammed in Figure 14(a). And should *porch* become well established in this usage, we can expect the

24. It is of course not an accident that some of these verbs are historically related to nouns and thus quite similar in form. Examples like *paint*, *wrap*, and *cover* are listed here, rather than in (15), based on the intuition that the verb sense is more basic.

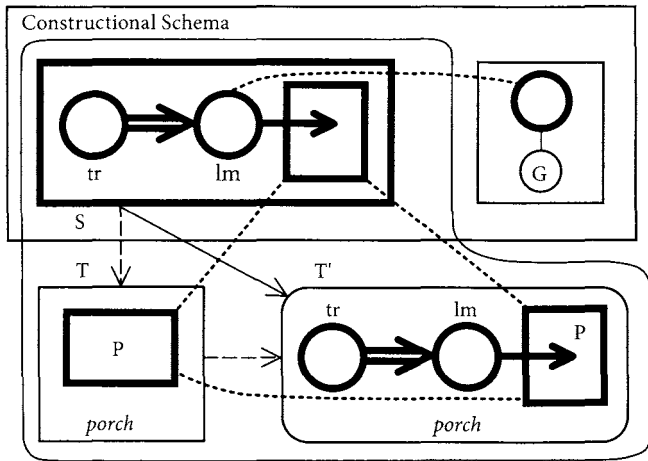


Figure 19. A case of derivation by construction

same kind of development shown for *kick* in Figures 14(b) and 15, leading eventually to a conventional verbal sense. Presumably this happened for the denominal verbs in (15)c. But there is more to this story. As things stand, the extension of *porch* to verbal use is treated as if it were an isolated, unprecedented occurrence. In fact, however, *porch* is following a well-trod path. While this particular extension may be novel, it instantiates a familiar pattern which itself constitutes a conventional unit of English.

In CG terms, this pattern is nothing more than a schematized version of Figure 19, representing the abstracted commonality of numerous instances of nouns being extended to verbal use in this manner. It is shown in Figure 20, which differs from 19 only in that the noun is schematic rather than specific. The entire configuration is an established unit – a symbolic assembly which incorporates a constructional schema as one substructure. We can describe it in different ways, all equivalent in the present perspective. From one standpoint it constitutes an augmentation of the constructional schema. It is also describable as a conventional usage pattern. It might equally well be characterized as a pattern of semantic extension or constructional derivation. However we describe it, the pattern is accessible for the sanction of novel usages. When *porch* is used as a verb, it instantiates the noun slot in Figure 20, resulting in Figure 19. Despite being lexicalized, the derived verbs in (15)c also reflect the pattern (just as lexical nouns like *teacher* and *computer* instantiate the productive pattern of *-er* derivation).

Here a subtle point arises concerning the distinction between a skewing usage and a skewing construction. When *porch* is assessed with respect to the constructional schema, as shown in Figure 19, it represents a skewing usage, since the noun is non-congruent with the schematic verb it instantiates. It is however congruent with the element it instantiates within the overall assembly in Figure 20, i.e. a schematic noun construable as a location. Assessed with respect to the overall pattern, therefore, the occurrence of *porch* is not a skewing usage but a straightforward case of elaboration. By its very nature,

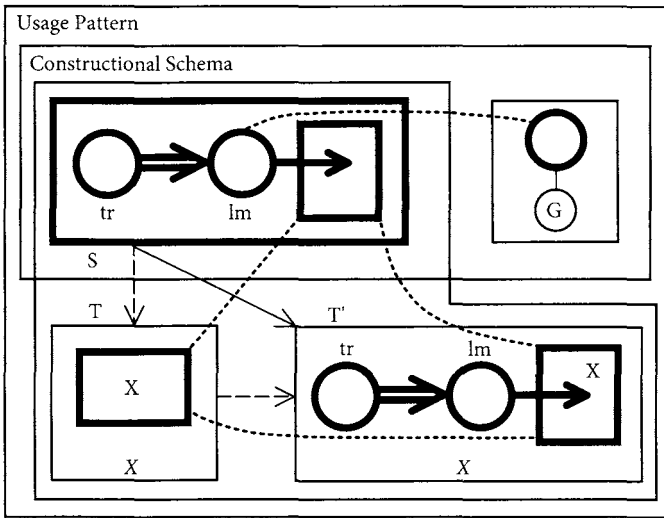


Figure 20. Schema for derivation by construction

though, this overall construction performs a derivational function and thus gives rise to new lexical senses.

Configurations analogous to Figure 20 can be posited for other verbalization patterns, like those reflected in (15)a–b. Since these are quite similar to one another, a more schematic assembly might well emerge that subsumes them.²⁵ One could also imagine a highly abstract assembly encompassing the use of nouns as verbs in other constructions (like the *way* construction in (14)c). Whatever the specific details, collectively these assemblies represent the noteworthy tolerance of English for extending nouns to verbal use.

8. Conclusion

I have attempted a unification at several levels. There is first the unification of lexicon, morphology, and syntax, all described in CG as assemblies of symbolic structures. Symbolic assemblies also permit a unified treatment of syntagmatic and paradigmatic relations, as well as derivation, semantic extension, and patterns of usage. More generally, symbolic assemblies are special cases of networks comprising structures linked by directional relationships. I have sketched a unified account of such relationships embracing a wide range of phenomena, including recognition, categorization, inheritance, composition, metaphor, metonymy, and blending. To be sure, the account is programmatic and in some respects speculative. I believe, however, that it makes sense linguistically and has some chance of proving cognitively realistic.

25. By the same token, they might be articulated into lower-level schemas for particular classes of nouns.

References

- Allwood, J. 2003. Meaning potentials and context: Some consequences for the analysis of variation in meaning. In *Cognitive Approaches to Lexical Semantics* [Cognitive Linguistics Research 23], H. Cuyckens, R. Dirven & J. R. Taylor (eds.), 29–65. Berlin: Mouton de Gruyter.
- Barlow, M. & Kemmer, S. (eds.). 2000. *Usage-Based Models of Language*. Stanford: CSLI Publications.
- Chomsky, N. 1965. *Aspects of the Theory of Syntax*. Cambridge MA: MIT Press.
- Clark, E. V. & Clark, H. H. 1979. When nouns surface as verbs. *Language* 55: 767–811.
- Croft, W. 2001. *Radical Construction Grammar: Syntactic Theory in Typological Perspective*. Oxford: Oxford University Press.
- Dowty, D. 2000. “The garden swarms with bees” and the fallacy of “argument alternation”. In *Polysemy: Theoretical and Computational Approaches*, Y. Ravin & C. Leacock (eds.), 111–128. Oxford: Oxford University Press.
- Fauconnier, G. 1985. *Mental Spaces: Aspects of Meaning Construction in Natural Language*. Cambridge MA: MIT Press.
- Fauconnier, G. 1997. *Mappings in Thought and Language*. Cambridge: Cambridge University Press.
- Fauconnier, G. & Sweetser, E. (eds.). 1996. *Spaces, Worlds, and Grammar*. Chicago: University of Chicago Press.
- Fauconnier, G. & Turner, M. 1998. Conceptual integration networks. *Cognitive Science* 22: 133–187.
- Fauconnier, G. & Turner, M. 2002. *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. New York: Basic Books.
- Fillmore, C. J. 1988. The mechanisms of “construction grammar”. *Proceedings of the Annual Meeting of the Berkeley Linguistics Society* 14: 35–55.
- Fillmore, C. J., Kay, P. & O'Connor, M. C. 1988. Regularity and idiomaticity in grammatical constructions: The case of *let alone*. *Language* 64: 501–538.
- Goldberg, A.E. 1995. *Constructions: A Construction Grammar Approach to Argument Structure*. Chicago: University of Chicago Press.
- Goldberg, A. E. & Jackendoff, R. 2004. The English resultative as a family of constructions. *Language* 80: 532–568.
- Haiman, J. & Ourn, N. 2003. Nouns, verbs, and syntactic backsliding in Khmer. *Studies in Language* 27: 505–528.
- Israel, M. 1996. The way constructions grow. In *Conceptual Structure, Discourse and Language*, A. E. Goldberg (ed.), 217–230. Stanford: CSLI Publications.
- Jackendoff, R. 1975. Morphological and semantic regularities in the lexicon. *Language* 51: 639–671.
- Jacobsen, W. H., Jr. 1979. Noun and verb in Nootkan. In *The Victoria Conference on Northwestern Languages* [Heritage Record 4], B. S. Efrat (ed.), 83–155. Victoria: British Columbia Provincial Museum.
- Kövecses, Z. & Radden, G. 1998. Metonymy: Developing a cognitive linguistic view. *Cognitive Linguistics* 9: 37–77.
- Lakoff, G. 1987. *Women, Fire, and Dangerous Things: What Categories Reveal About the Mind*. Chicago: University of Chicago Press.
- Lakoff, G. & Johnson, M. 1980. *Metaphors We Live By*. Chicago: University of Chicago Press.
- Langacker, R. W. 1982. Space grammar, analysability, and the English passive. *Language* 58: 22–80.
- Langacker, R. W. 1987. *Foundations of Cognitive Grammar*, Volume 1, *Theoretical Prerequisites*. Stanford: Stanford University Press.
- Langacker, R. W. 1991. *Foundations of Cognitive Grammar*, Volume 2, *Descriptive Application*. Stanford: Stanford University Press.
- Langacker, R. W. 1999. Double-subject constructions. In *Linguistics in the Morning Calm* 4, S.-Y. Bak (ed.), 83–104. Seoul: Hanshin.
- Langacker, R. W. 2000. A dynamic usage-based model. In *Usage-Based Models of Language*, M. Barlow & S. Kemmer (eds.), 1–63. Stanford: CSLI Publications.

- Langacker, R. W. 2003. Constructions in cognitive grammar. *English Linguistics* 20: 41–83.
- Langacker, R. W. 2004. Remarks on nominal grounding. *Functions of Language* 11: 77–113.
- Langacker, R. W. 2005a. Integration, grammaticization, and constructional meaning. In *Grammatical Constructions: Back to the Roots* [Constructional Approaches to Language 4], M. Fried & H. C. Boas (eds.), 157–189. Amsterdam: John Benjamins.
- Langacker, R. W. 2005b. Construction grammars: Cognitive, radical, and less so. In *Cognitive Linguistics: Internal Dynamics and Interdisciplinary Interaction* [Cognitive Linguistics Research 32], F. J. Ruiz de Mendoza Ibáñez & M. S. Peña Cervel (eds.), 101–159. Berlin: Mouton de Gruyter.
- Langacker, R. W. 2006. On the continuous debate about discreteness. *Cognitive Linguistics* 17: 107–151.
- Michaelis, L. A. 2004. Type shifting in construction grammar: An integrated approach to aspectual coercion. *Cognitive Linguistics* 15: 1–67.
- Michaelis, L. A. & Lambrecht, K. 1996. Toward a construction-based theory of language function: The case of nominal extraposition. *Language* 72: 215–247.
- Panther, K.-U. & Radden, G. (eds.). 2004. *Metonymy in Language and Thought* [Human Cognitive Processing 4]. Amsterdam: John Benjamins.
- Queller, K. 2001. A usage-based approach to modeling and teaching the phrasal lexicon. In *Applied Cognitive Linguistics II: Language Pedagogy* [Cognitive Linguistics Research 19.2], M. Pütz, S. Niemeier & R. Dirven (eds.), 55–83. Berlin: Mouton de Gruyter.
- Queller, K. 2003. Metonymic sense shift: Its origins in hearers' abductive construal of usage in context. In *Cognitive Approaches to Lexical Semantics* [Cognitive Linguistics Research 23], H. Cuyckens, R. Dirven & J. R. Taylor (eds.), 211–241. Berlin: Mouton de Gruyter.
- Sandra, D. & Rice, S. 1995. Network analyses of prepositional meaning: Mirroring whose mind – the linguist's or the language user's? *Cognitive Linguistics* 6: 89–130.
- Tomasello, M. 2003. *Constructing a Language: A Usage-Based Theory of Language Acquisition*. Cambridge MA: Harvard University Press.
- Turner, M. 1987. *Death is the Mother of Beauty: Mind, Metaphor, Criticism*. Chicago: University of Chicago Press.
- Zlatev, J. 2003. Polysemy or generality? Mu. In *Cognitive Approaches to Lexical Semantics* [Cognitive Linguistics Research 23], H. Cuyckens, R. Dirven & J. R. Taylor (eds.), 447–494. Berlin: Mouton de Gruyter.

Partonomic structures in syntax

Edith A. Moravcsik

1. The utility of positing partonomic structures

How do linguists formulate terms for grammatical rules? For an example, consider a rule of syntax:

- (1) The adjective precedes the noun.

This rule has two logical components: the predicate “precedes”, and the terms “adjective” and “noun” for which this predicate is said to hold. The conceptual tools involved in creating these terms are *segmentation* and *classification* (or categorization). Starting with sentences as wholes, the labels adjective and noun are based on the assumption that sentences can be segmented into words, and that some words within and across sentences are alike in some ways so that they can be placed into categories. Term formation is thus based on invoking the syntagmatic relation of *partonomy* (whole-part relations) and the paradigmatic relation of *taxonomy* (type-token relations).

In this paper, the rationale for the first of these two relations will be discussed: what kinds of partonomic relations are posited in syntax and why? The paper is a study in cognitive metalinguistics. “Cognitive” because it has to do with interpretations that the human mind imposes on reality; and pertaining to metalinguistics because it is about interpretations that the analyst imposes on language rather than those that speakers of a language impose on extralinguistic reality (on the latter, see for example Moltmann 1997).¹

Why is it useful to posit partonomic relations? The short answer is that it facilitates generalizations. Partonomic interpretation can proceed in two directions: either an entity is assumed as given and we posit parts within it, or a set of entities is assumed to be given and we posit a whole that subsumes them as parts. We will label the two conceptual moves as analysis and synthesis, respectively. Let us now see how these concepts facilitate generalizations.

Analysis – breaking entities into parts – allows us to pinpoint similarities among entities that would otherwise seem different.

1. On the still young and developing field of cognitive metalinguistics and, more broadly, the cognitive science of science, see Grier (1988), Kertész (2004a, 2004b), Kertész and Rákosi (2005), and Rákosi (2005).

- (2) The utility of partonomic analysis: Different wholes may have similar parts.

This is so in all domains of inquiry. For example, two different animal groups, such as birds and insects, are similar in that both may have wings; and different substances can be brought to common denominators by chemists and physicists if they are analyzed into smaller and smaller elementary parts. In phonology, linguists segment words into sounds for the same reason; for example, the words *crave* and *spin* are similar in that both start with a consonant cluster. And in syntax, the two sentences *Should I call you?* and *Never has Jack slept better* turn out to be similar when we break them into words: both contain the auxiliary before the subject.

Recognizing parts of a whole may also illuminate the nature of that whole. A whole may be well-formed or ill-formed due to its parts. This is borne out by auto mechanics identifying a part of a broken-down car as responsible for the problem and by doctors looking at a sick body and finding the part that causes the disease. Similarly in syntax: the ungrammaticality of the sentence *Jack did slept better* is due to the joint occurrence of the parts *did* and *slept*. The assumption behind such analyses is compositionality: the nature of the whole is determined in some way by the nature of the parts and their relationships.

These examples illustrate the utility of breaking larger entities into smaller ones. As mentioned above, partonomic analysis may alternatively start with a set of entities which are then synthesized into a single whole. The motivation for doing so is that even though the components are different, the wholes they form may be similar.

- (3) The utility of partonomic synthesis: Different parts may form similar wholes.

This is the recognition that astronomers appear to act on when they group certain sets of diverse heavenly bodies into solar systems; when different sets of pathological symptoms are identified as pointing to the same disease; when phonologists recognize groups of different sounds as forming syllables; and when syntacticians say that different words may form the same type of phrase. For example, in the sentences *The new employees are quitting* and *They have escaped*, the different words *the new employee* and *they* form the same type of whole: a plural subject.

All in all, in partonomic analysis, larger things are analyzed into parts or smaller things are synthesized into wholes. Either way, partonomic analysis legitimizes what would otherwise be a paradox: one thing declared to be several things and several things declared as one thing. “One” and “more than one” are contradictory notions: on the face of it, something cannot be both one and many. But partonomic analysis allows us to re-state the paradox so that it is not contradictory any more: “one” can be “more-than-one” and “more-than-one” can be “one” if a whole consists of more than one part.

In Sections 2–5, we will discuss some examples of partonomy as a problem-solving tool in syntax. In Section 6, similar uses of partonomy in other sciences and in everyday thought will be cited.

2. Synthesis: Positing wholes for parts

Syntacticians posit wholes for parts when they hypothesize that words form phrases, clauses, sentences, paragraphs, and discourses.² In what follows, we will focus on phrases and clauses.

Arguments for phrases and clauses are known as “constituency tests” (cf. for example Radford 1981:34–78; McCawley 1988:55–66; Croft 2001:185–195). Although the term “constituency test” suggests that the goal is to establish constituents, the ultimate goal is to state generalizations; constituents are of interest only to the extent that they support rules. Thus, each constituent test is actually an appeal to a generalization that is facilitated by the assumption of a constituent.

The point of constituent tests is that certain assemblages of words act in concert: they act as if they were one. This “acting like one” is borne out in two basic ways: dependence among the words of the phrase and independence of the words of one phrase from the words of other phrases. Internal dependence within a syntactic phrase and external independence among syntactic phrases are manifested in various ways as shown in (4).³

- (4) a. Internal dependence:
 (i) joint recurrence within and across sentences;
 (ii) joint non-occurrence (through replacement or omission) in a sentence;
 (iii) contiguity.
 b. External independence: joint occurrence as a sentence

These criteria and how the constituents they define facilitate syntactic rules are illustrated for the noun phrase in (5)–(9).

- (5) Joint recurrence within and across sentences
 Subject and object
 a. Rule with noun phrase not assumed:
 “What may serve as subject and object is *article and adjective and noun*.”
 b. Rule with noun phrase assumed:
 “What may serve as subject and object is *a noun phrase*.”

For example:

[*The new students*]_{NP} bought [*an interesting book*]_{NP}.
 [*A large spider*]_{NP} was crawling on [*the newly-painted wall*]_{NP}.

2. For discussions of constituent structure in syntax and how they emerge in the course of language use, see Jakobson (1963), Lakoff (1987:283–285), Speas (ed.) (1990), Leffell and Bouchard (ed.) (1991), Langacker (1997, 1999), Bybee and Scheibman (1999), and the papers in Bybee and Hopper (ed.) (2001), especially those on pages 1–24, 229–428, and 449–470. For a brief overview, see Moravcsik, to appear. For the usefulness of partonomic analysis in understanding historical change in language and culture, see Enfield (2005:194–197).

3. In addition to syntactic criteria for phrasehood, there are also semantic, morphological, and phonological ones, the last including words forming a single stress or pitch group and manifesting phonological interaction other than those that occur across phrases. On the role of prosodic cues for grouping, see Hunyadi, to appear.

- (6) Joint non-occurrence through replacement
- a. Rule with noun phrase not assumed:
“What may be replaced by a pronoun is *article and adjective and noun*.”
 - b. Rule with noun phrase assumed:
“What may be replaced by a pronoun is *a noun phrase*.”

For example:

[*The out-of-town guests*]_{NP} arrived and *they* left immediately.

- (7) Joint non-occurrence through omission
- a. Rule with noun phrase not assumed:
“What may be omitted under referential identity is *article and adjective and noun*.”
 - b. Rule with noun phrase assumed:
“What may be omitted under referential identity is *a noun phrase*.”

For example:

[*The out-of-town guests*]_{NP} arrived and ___ left immediately.

- (8) Contiguity
- a. Rule with noun phrase not assumed:
“Words that must be contiguous are *article and adjective and noun*.”
 - b. Rule with noun phrase assumed:
“Words that must be contiguous are those belonging to *a noun phrase*.”

For example:

[*The out-of-town guests*]_{NP} arrived.

**The arrived out-of-town guests*.

- (9) Joint occurrence as a sentence
- a. Rule with noun phrase not assumed:
“What may be an answer sentence to a question is *article and adjective and noun*.”
 - b. Rule with noun phrase assumed:
“What may be an answer sentence to a question is *a noun phrase*.”

For example:

Question: *Who arrived last night?*

Answer: [*The out-of-town guests*]_{NP}.

3. Problems with synthesis

3.1 Complexity


While, as seen above, the assumption of phrases facilitates generalizations, it can also create problems. Two of these problems are complexity and inconsistency: phrases may turn out to be *complex*; and evidence for phrasehood may be *contradictory*.

Let us first consider complexity. A partonomy is simple if it involves minimal structure. In the diagrams of (10), WH stands for ‘whole’, P1, P2 etc. stand for parts; lines trace partonomic relations; complexities are crossed out.

(10) Minimal partonomic structure:


(A) ONLY TWO PARTS

Each whole contains only two sister-parts.




(B) ONLY TWO LEVELS

Parts do not contain further parts.




(C) PARTS ARE EQUAL

Sister-parts are of equal rank.




(D) PARTS ARE UNIQUELY ASSIGNED TO WHOLES

Every part belongs to one and only one directly superordinate whole.



(E) THE WHOLE IS COMPOSITIONAL

The whole is compositional: its characteristics are predictable from the characteristics of the parts and their relations.



Actual partonomic structures in syntax often deviate from one or more of these desiderata of simplicity. Here follow some examples of complex partonomies.

(11) Complex partonomic structures

(A) MORE THAN TWO PARTS

Example: *John, Mary, and Sue*

(B) MORE THAN TWO LEVELS

Example: *[two old [japanese cars]] and [a new [one]]*

(C) PARTS ARE NOT EQUAL

Example: *walk fast*

(D) PARTS ARE NON-UNIQUELY ASSIGNED TO WHOLE

Example: *I expect him to leave town.*

(E) THE WHOLE IS NOT COMPOSITIONAL

Example: *Mary was born in London and Bill, in Chicago.*


For example, while in *He left town* (10D), each word belongs to only one part of the sentence (a noun phrase, a verb, and a second noun phrase), in *I expect him to leave town* (11D), the word *him* seems to belong both to the main clause and to the subordinate clause. In (10E), *Bill was born in Chicago*, the meaning of the sentence is the sum of the meanings of the parts and their relations but in (11E) *Mary was born in London and Bill, in Chicago*, the verbal meaning 'was born' has no overt expression in the second clause.

Let us expand on one of the best-documented types of complexity in partonomic structure: the asymmetry of parts (C) in (11). Other than in coordinate structure, the parts of a syntactic phrase are never fully equal by token of the very fact that they bear different category labels, such as Adjective and Noun, or Verb and Adverb. But beyond this, there are also cross-categorial asymmetries within phrases: for example, nouns of noun phrases and verbs of verb phrases show similar behavior even though they belong to different categories: noun and verb (cf. Corbett et al. 1993; Croft 2001: 241–280).

Such cross-categorial asymmetries are borne out both by the selection and the ordering of syntactic constituents. First, given a phrase, one of its components may be able to stand by itself in the same context where the phrase occurs but other components cannot. Thus, in the noun phrase *brown dogs*, the noun is syntactically obligatory but the adjective is optional; and in verb phrases such as *run fast*, the verb is obligatory but the adverb is dispensable. Second, there is a crosslinguistic tendency for a language to morphologically mark (such as for case or agreement) either heads or dependants across different kinds of phrases (Nichols 1986). Third, languages show some tendency toward the uniform ordering of different kinds of constituents whose classes cut across word and phrase types. This fact has been interpreted differently, with the division being heads versus dependants (Vennemann 1973), or branching versus non-branching constituents (Dryer 1992), or mother-node-constructing versus non-mother-node-constructing constituents (Hawkins 1994). But whatever the common denominators posited for constituents that tend to be uniformly ordered in any one language, they highlight an asymmetry within phrases.

The recognition of the non-equality of phrase structure has given rise to dependency grammars which explicitly represent this asymmetry. In Richard Hudson's word grammar (1984, 1990), dependency relations among words are primary and the concept of wholes – phrases – is derivative. In this framework, the whole is implicit, defined as a head along with its dependants. The difference between the two representations is shown in (12).

(12) Phrases in constituent structure and in dependency grammar



The examples seen so far (cf. (11)) illustrated one of the two types of problems that arise when wholes are posited: complexity. A second problem that arises is contradictory evidence, which we turn to next.

3.2 Inconsistency

A partonomic structure is consistent if its parts act as a whole by all the relevant criteria listed in (4) above. An example is *running races*: as shown in (13), this is a phrase by all four criteria.

- (13) a. joint recurrence in and across sentences
 Example: *Running races is exciting but Joel dislikes running races.*
Running races is difficult.
- b. joint non-occurrence through replacement or omission
 Example: *Running races is exciting but Joel dislikes it.*
Running races is exciting but ___ difficult.
- c. contiguity
 Example: *Running races is exciting.*
 **Running is exciting races.*
- d. occurrence as a sentence
 Example: Question: *What is Joel's favorite activity?*
 Answer: *Running races.*

In contrast, a partonomic structure will be said to be inconsistent if its parts act as a whole by some criteria but not by others.⁴ This is frequently so in syntax (cf. Croft 2001: 185–197). Here are two examples.

- (14) a. A set of words makes a phrase by joint replaceability but not by contiguity:
The man is a friend of yours who came to see me.
He is from Chicago.
 **He is a friend of yours who came to see me.*
- b. A set of words makes a phrase by joint replaceability but cannot stand as a sentence:
Bill bought an old Japanese car and Jill bought a new one.
 Question: *What did Jill buy?*
 Answer: **Japanese car.*

In the two parts of Section 3, we have seen instances of complex and inconsistent partonomies. Let us now turn to the question of how they are addressed in syntactic description.


4. The concepts of complexity and inconsistency are not fully distinct: complexities may be viewed as inconsistencies relative to the requirement of simple partonomies.

4. Solving problems of synthesis with analysis: Positing parts for wholes

4.1 Eliminating complexity

As noted above (cf. (11) (D)), one example of complex partonomic structures involves overlapping constituents. The problem is schematized in (15).

(15) Overlapping constituents




For example, in *I expect him to leave home*, P3 is *him*.

One solution involves teasing out two separate levels, or layers, of this structure so that each is free of overlapping constituency: on each level, every part belongs to only one whole. We will term this kind of partonomic analysis *layering*.

Layering is based on the idea that things can be viewed from different angles and depending on the point of view, their structure may be different. Partonomic structure appears complex only as long as the different aspects of the construction are not separated; when they are, each of the resulting structures is simple. Layering is schematized in (16), where each boxed structure represents a separate layer.

(16) Eliminating overlapping constituents by layering



The layering of complex syntactic structures has taken various forms in the literature (cf. Moravcsik 2006, especially Chapters 1 and 2). Minimal layering involves a partonomic diagram with *two faces*. It is illustrated by the example of long-distance verb-object agreement in Hungarian (É. Kiss 1987: 224–273).

In this language, transitive verbs agree not only with the subject (in person and number) but also with the direct object (in definiteness and person). Consider the main verb in (17).

- (17) *Én szeretné-lek látni téged.*
 I would:like-S1_s:S2_O to:see you_s
 'I would like to see you.'

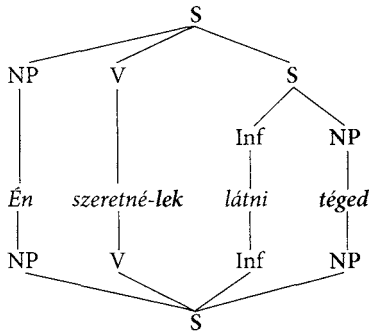
First, it shows agreement in number and person with the subject as expected. Second, it also shows object agreement but there is a problem here: the main verb agrees not with its own object – which would be the clause ‘to see you’ – but with the object of that clause:

'you'. Agreement is thus non-local, or long-distance: agreement controller and agreement target are not clause-mates as they usually are.

This agreement pattern indicates overlapping constituency: the main clause and the subordinate clause overlap in the subordinate object 'you'. On the one hand, this constituent is part of the subordinate clause as shown by the fact that it is selected and case-marked by the subordinate verb and ordered relative to it. On the other hand, this constituent is also part of the main clause in that the main verb agrees with it and in that it may alternatively be ordered into focus position preceding the main verb, as in *Én téged szeretnélek látni*.

Considering the contradictory evidence, É. Kiss proposes a bifacial tree. The lower face shows the sentence as biclausal with the object belonging to the subordinate clause. The upper face shows the sentence as monoclausal, thus legitimizing the agreement of the verb with the object, with locality of agreement upheld. A simplified form of the proposed tree structure is shown in (18).

(18) Layering into two faces of a structure



In this account, a single sentence structure is layered into two structures, both of which are simultaneously present in the syntactic derivation. Other varieties of layering assume multiple levels where one level is an input to a rule creating another level. In some accounts, the levels are in the same grammatical component.

The first kind is documented in the long history of transformational generative grammar where different levels of syntactic structure have been assumed. For an example of multiple syntactic levels resolving complex partonomies, consider another instance of overlapping constituency known as raising constructions. The sentence mentioned earlier – *I expect him to leave home* – is an example and so is (19).

(19) *Bill expects her to pass the exam.*

Just as in the Hungarian case, a noun phrase seems to be part of both the main clause and the subordinate clause. Evidence for *her* being part of the subordinate clause is the selectional relations it bears to the subordinate verb. Evidence for it belonging to the main clause is case marking: the objective case of *her* is assigned by the main verb *expects*. The solution consists in the sentence representation layered into two levels: deep structure and surface structure. On each level, the noun phrase is part of only one clause (Postal 1974).

- (20) Layering into two syntactic structures

Surface structure:

[Bill expects her to pass the exam.]_S

Underlying structure:

[Bill expects [that she passes the exam.]]_S

The above examples show how layering a complex structure into two simple structures can resolve some *complexities* of partonomic structure – in particular, that of overlapping parts. Of the two anomalies noted in Section 3 – complexity and inconsistency – complexity is less of a problem: the scientific criterion that it violates is simplicity, which is basically an esthetic criterion. The larger problem is inconsistency since it conflicts with the requirement that scientific accounts be free of contradictions.


The same idea of layering problematic structures into two non-problematic ones that has been evoked to resolve complexities has been resorted to for resolving *inconsistencies* in constituent structures. A prime example is discontinuous constituents.

4.2 Eliminating inconsistency


The term discontinuous constituency refers to an assemblage of words that form a constituent by some criteria but not by the criterion of contiguity (Huck and Ojeda (ed.) 1987; Bunt and Horck (ed.) 1996; Croft 2001: 186–188). An example was seen in (14) (a) above: in the sentence *The man is a friend of yours who came to see me*, the words *the man ... who came to see me* form a whole by the criterion of replacement but they are discontinuous.

Discontinuity has been a central problem in generative grammar, with the solution being distinct levels of syntactic representation connected by movement rules. An example of discontinuity and its solution by layering is schematized in (21).

- (21) a. Discontinuity



- b. Eliminating discontinuity by layering



For example, the sentence *The letter came that you were waiting for* can be analyzed as in (22).

- (22) Layering into two syntactic structures
 Surface structure:
The letter came that you were waiting for.
 Underlying structure:
The letter that you were waiting for came.

5. Layering in cognitive grammar

In approaches that assume multiple syntactic representation to eliminate complex or inconsistent partonomic structures, there has been a continuing concern to motivate the two syntactic levels into which such structures are sliced. Several mechanisms were designed to address this issue.

One type of attempt has been aimed at *reducing the differences* between the two levels. This is the thrust of Joseph Emonds' Structure Preserving Constraint; of Noam Chomsky's Projection Principle; and of the requirement posited in some versions of generative grammar that derivations be monotonic – i.e. allowing additions to a structure but not diminishing it or replacing it.

The other type of attempt has been to *independently motivate* the two structures. Multiple syntactic levels have been recognized in several frameworks – including the Minimalist Program proposed by Noam Chomsky – as arbitrary. A layering that is, however, multiply-motivated is one where the layers are meaning and syntactic form. In Sadock's Autolexical Grammar (1991) and in Langacker's Cognitive Grammar (e.g. 1997, 1999), discontinuity is accounted for by slicing sentence representations into a meaning level and a form level; or, in Langacker's framework, into a conceptual level and a phonological level.

- (23) Layering into form structure and meaning structure:
 Form:
The letter came that you were waiting for.
 Meaning:
The letter that you were waiting for came.
 (Langacker 1997: 25)

In these frameworks, overlapping phrases and discontinuous phrases are re-analyzed as regular rather than exceptional: their apparent complexity or inconsistency is the result of the analyst failing to discern two distinct aspects of them: form and meaning. Mismatches between these two are expected since meaning and form are independently known to be distinct kinds of entities. The same conceptual distinction is seminal in Croft's Radical Cognitive Grammar as well: semantics and syntax are seen as independent entities (e.g. Croft 2001: 108).

In such approaches, the conflict of a single entity – syntactic structure – being complex or outright self-contradictory simply does not arise. In fact, the notion of syntactic constituent structure as an autonomous layer is explicitly eliminated by Langacker (2005: 103–112). In his framework, partonomic structures exist only in pronunciation and in conceptualization. Whether phonological and conceptual partonomies are entirely free of complexities and inconsistencies of the sort surveyed above remains to be seen as studies in cognitive grammar continue to evolve.

6. Partonomy as a ubiquitous cognitive tool

As noted in Section 1, the paradoxical relationship between ‘one’ and ‘more than one’ has been a central issue not only in linguistics but in science in general, in philosophy, and even in theology. Are the Father, the Son, and the Holy Ghost one or more than one? The concept of the Trinity, first conceived by Tertullian around 200 AD, has been accepted by some theologians as consistent with monotheism while rejected by others as a sign of leaning towards polytheism. Tertullian’s own analysis is “*tres personae, una substantia*”, which may perhaps be interpreted as invoking partonomic structure for solving the paradox: three persons being parts of the single substance.

Some examples of partonomic structure in science have already been cited in Section 1.⁵ The most prominent use of partonomy in science is analysis of wholes into parts – a basic tool in physics, biology, sociology, and other fields. Although segmenting things into parts is a dominant methodology in science, the opposite also occurs: positing larger wholes for sets of individual objects. As mentioned earlier, this is what happens when individual astronomical objects are subsumed under larger wholes such as solar systems and galaxies; or when individual symptoms are identified as parts of a single disease.

Analysis of wholes into parts is also fundamental to ordinary people’s perception and interpretation of the world. It is present when a child takes apart a toy car. The relation between a whole and a part is known by children very early – earlier than the somewhat parallel relation between type and token (Markman 1989: 161–233). For example, children learn that oak is part of the forest before they learn that it is a subtype of trees.

The reverse – creating wholes for parts, called chunking in the psychological literature (cf. MacWhinney 2005: 91–92) – is similarly ubiquitous as a general human tool of dealing with complexity. Individual digits of telephone numbers, credit card numbers, and social security numbers are commonly re-interpreted as sequences of larger units. Gleitman (1981) shows how chunking is an aid in verbal memory and in problem solving (288–289, 319–324). In their classic study, William Bryan and Noble Harter (1899) demonstrated how the acquisition of the skills of a telegrapher decoding incoming messages was based

5. For general analyses of part-whole relations and how they figure in scientific and everyday arguments, see Lerner (ed.) (1963), Husserl (1970) (especially Volume II, Investigation III), Winston, Chaffin and Herrmann (1987), Kertész and Rákosi (2005), Varzi no date, and Burkhardt, Seibt and Imaguire (eds) (to appear).

on what they called “a hierarchy of habits”: first learning to decode syllables, then words, then phrases, then sentences. Herbert Simon provides an evolutionary explanation for the hierarchic – part-whole – structure of the world and/or for the fact that humans analyze it in this manner (Simon 1996: 183–216).

The criteria that are used for wholes in science and in everyday thinking – internal dependence and external independence – are similar to those used in syntactic description. As we saw above, one criterion for parts forming a whole is contiguity. The general expectation that parts of a whole be adjacent is borne out in human perception in general. For example, Max Wertheimer (1938) points out that one factor involved in the natural grouping of both visual and auditory stimuli is proximity: people tend to unite things that are proximate more than those that are not. He calls it the Law of Proximity (cf. also Hunyadi, to appear). Another, related principle of Wertheimer’s is “the Factor of Uniform Destiny”: elements that cohere are shifted together. This criterion evokes generative grammar’s criterion of joint movability for constituents. The zoologist Jacob von Uexhuell (1921:7) similarly states: “Ein Gegenstand ist was sich zusammen bewegt.” (‘An object is made up of whatever moves together.’)

Let us now turn to partonomic complexities and inconsistencies. Complex and inconsistent partonomic patterns are frequent in scientific thinking and in everyday thought. Two examples of rampant complexity are the asymmetry of parts and the lack of compositionality.

Things are commonly perceived as having *asymmetric parts*. As Richard Hudson has pointed out in some of his writings (e.g. 1984: 38) and has been amply documented by Barbara Tversky (1990), people perceive wholes as consisting of more representative and less representative components, such as house and garage, or head and trunk.

Scientists, just as syntacticians, struggle with complex partonomic structures, such as *lack of compositionality*. In systems thinking, it is recognized that wholes may be more than the sum of their parts and relations (cf. Hookway 2000). Physicist Fritjof Capra provides a concrete illustration of non-compositionality in the physical world in the following passage (Capra 1996: 28–29):

At each level of complexity, the observed phenomena exhibit properties that do not exist at the lower level. For example, the concept of temperature, which is central to thermodynamics, is meaningless at the level of individual atoms, where the laws of quantum theory operate. Similarly, the taste of sugar is not present in the carbon, hydrogen, and oxygen atoms that constitute its components. In the early 1920s, the philosopher C. D. Broad coined the term ‘emergent properties’ for those properties that emerge at a certain level of complexity but do not exist at lower levels.

In everyday thinking, too, the concept of non-compositionality is a recurrent theme; the phrase “the whole is more than the sum of its parts” is frequently used in everyday parlance.

As we saw above, partonomic analysis – splitting things into parts – can be a tool of conflict resolution in syntactic description. It has a similar role in how we come to terms with the world’s complexities and contradictions in general. This is borne out in the ways people and societies deal with conflict: countries that have internal incongruities are split

into two and so are political parties, dissenting religions, and conflicted marriages. In each case, conflict within a single entity is solved by splitting that entity so that each of the resulting units is itself free of conflicts (cf. Husserl 1970: 754–759).

How partonomic analysis is used in everyday thinking for resolving contradictions is well illustrated in a study by Sharpe, Eakin, Saragovi and Macnamara 1996. The purpose of this research was to see how people come to grips with contradictions. The subjects were 40 undergraduates; they were instructed to provide a free response to a question in half a page or less. The question had to do with the following situation. A student asks her professor whether her term paper is good. The professor pauses and says “Yes and no” – a contradictory response. The question posed to the subjects was this: “Can you make sense of this answer? If so, please explain how.” The results showed that appealing to partonomic structure was the dominant strategy: 97.5% of the subjects said that part of the paper must have been good and part of it bad. A similar experiment with 24 children between 3;1 and 4;2 showed similar results. This study illustrates that people resolve a contradiction by splitting the conflicted object into two parts, each internally consistent.

All in all, partonomic analysis is a basic conceptual tool in how people perceive and interpret the world and it forms a common link between linguistics, general scientific theorizing and everyday human thinking.

7. Conclusions

This paper argued that partonomic analysis – the assumption of wholes and parts – is a useful device in syntactic description for two reasons. First, “vertical” partonomy – building wholes out of the words of a sentence – facilitates generalizations but often at the price of creating complexities and inconsistencies. Second, another application of partonomy, layering – i.e. “horizontal” slicing – helps accommodate these problems. As was seen, complexities such as overlapping constituents may be solved by layering sentence representations into two simultaneous faces or into representations on different levels in syntax or in different components. Similarly, partonomic inconsistencies such as discontinuous constituency may be solved by assuming different levels of representation. The case studies revealed that several of the well-known central theoretical constructs posited in syntactic description boil down to various applications of partonomy.

Furthermore, we have seen that partonomic synthesis and analysis are ubiquitous conflict-resolving conceptual tools in science and in everyday thinking.

Contemplating the essence of categorization, Cecil Brown defines a category as resulting from “the treatment of two or more distinguishable entities *as if they were the same*” (emphasis added; Brown 1990: 17). An analogous description holds for creating partonomies: it is treating two or more entities *as if they were one*; or, conversely, treating a single entity *as if it were more than one*. Both taxonomy (categorization) and partonomy legitimize what would otherwise appear to be self-contradictory notions. Taxonomy tackles an apparent qualitative inconsistency: that two things can be both the same and different.

Partonomy tackles a quantitative conflict: that two things can be both many and one.⁶ Both relations thus serve as crucial conceptual tools in interpreting the world, including, as this paper attempted to show for partonomy, syntacticians interpreting language structure.

Acknowledgements

For their very helpful comments, I am grateful to Csilla Rákosi, Vyv Evans, a reviewer of this volume, and the audience at the New Directions in Cognitive Linguistics Conference.

References

- Brown, C. H. 1990. A survey of category types in natural language. In *Meanings and Prototypes*, S. L. Tsohatzidis (ed.), 17–47. London: Routledge.
- Bryan, W. L. & Harter, N. 1899. Studies on the telegraphic language. The acquisition of a hierarchy of habits. *The Psychological Review* VI (4): 345–375.
- Bunt, H. & van Horck, A. (eds). 1996. *Discontinuous Constituency*. Berlin: Mouton de Gruyter.
- Burkhardt, H., Seibt, J. & Imaguire, G. (eds). To appear. *The Handbook of Mereology*.
- Bybee, J. & Hopper, P. (eds). 2001. *Frequency and the Emergence of Linguistic Structure*. Amsterdam: John Benjamins.
- Bybee, J. & Scheibman, J. 1999. The effect of usage on degrees of constituency: The reduction of *don't* in English. *Linguistics* 37 (4): 575–596.
- Capra, F. 1996. *The Web of Life*. New York: Doubleday.
- Corbett, G. G., Fraser, N. M. & McGlashan, S. (eds). 1993. *Heads in Grammatical Theory*. Cambridge: Cambridge University Press.
- Croft, W. 2001. *Radical Construction Grammar. Syntactic Theory in Typological Perspective*. Oxford: Oxford University Press.
- Dryer, M. 1992. The Greenbergian word order correlations. *Language* 68: 81–138.
- É. Kiss, K. 1987. *Configurationality in Hungarian*. Budapest: Akadémiai Kiadó.
- Enfield, N. J. 2005. Areal linguistics and Mainland Southeast Asia. *Annual Review of Anthropology* 34: 181–206.
- Gleitman, H. 1981. *Psychology*. New York: Norton.
- Grier, R. N. 1988. *Explaining Science. A Cognitive Approach*. Chicago: University of Chicago Press.
- Hawkins, J. A. 1994. *A Performance Theory of Order and Constituency*. Cambridge: Cambridge University Press.
- Hookway, C. 2000. Holism. In *A Companion to the Philosophy of Science*, W. H. Newton-Smith (ed.), 162–164. Oxford: Blackwell.
- Huck, G. J. & Ojeda, A. E. (eds). 1987. *Discontinuous Constituency. Syntax and Semantics*, Volume 20. Orlando: Academic Press.
- Hudson, R. 1984. *Word Grammar*. Oxford: Blackwell.
- Hudson, R. 1990. *English Word Grammar*. Oxford: Blackwell.
- Hunyadi, L. To appear. Grouping, the cognitive basis of recursion in language.
- Husserl, E. 1970. *Logical Investigations*. New York: Humanities Press.

6. For parallels between taxonomy and partonomy, see Tversky (1990).

- Jakobson, R. 1963. Parts and wholes in language. In *Parts and Wholes, The Hayden Colloquium on Scientific Method and Concept*, D. Lerner (ed.) 157–162. New York: The Free Press of Glencoe.
- Kertész, A. 2004a. *Cognitive Semantics and Scientific Knowledge. Case Studies in the Cognitive Science of Science*. Amsterdam: John Benjamins.
- Kertész, A. 2004b. *Philosophie der Linguistik. Studien zur naturalisierten Wissenschaftstheorie*. Tübingen: Gunter Narr.
- Kertész, A. & Rákosi, C. 2005. Whole-part and part-whole inferences in generative and cognitive linguistics. *Acta Linguistica* 52 (2–3): 221–280.
- Lakoff, G. 1987. *Women, Fire, and Dangerous Things. What Categories Reveal about the Mind*. Chicago: University of Chicago Press.
- Langacker, R. W. 1997. Constituency, dependency, and conceptual grouping. *Cognitive Linguistics* 8 (1): 1–32.
- Langacker, R. W. 1999. Conceptual grouping and constituency. In *Grammar and Conceptualization*, R. W. Langacker (ed.), 147–170. Berlin: Mouton de Gruyter.
- Langacker, R. W. 2005. Construction Grammars: cognitive, radical, and less so. In *Cognitive Linguistics: Internal Dynamics and Interdisciplinary Interaction*, F. J. Ruiz de Mendoza Ibanez & M. Sandra Pena Cervel (eds), 101–159. Berlin: Mouton de Gruyter.
- Leffel, K. & Bouchard, D. (eds). 1991. *Views on Phrase Structure*. Dordrecht: Kluwer.
- Lerner, D. (ed.). 1963. *Parts and Wholes, The Hayden Colloquium on Scientific Method and Concept*. New York: The Free Press of Glencoe.
- MacWhinney, B. 2005. New directions in the competition model. In *Beyond Nature-Nurture. Essays in honor of Elizabeth Bates*, M. Tomasello & D. I. Slobin (eds), 81–110. Mahwah, NJ: Erlbaum.
- Markman, E. M. 1989. *Categorization and Naming in Children. Problems of Induction*. Cambridge, MA: MIT Press.
- McCawley, J. D. 1988. *The Syntactic Phenomena of English*. Volume I. Chicago: University of Chicago Press.
- Moltmann, F. 1997. *Parts and Wholes in Semantics*. New York: Oxford University Press.
- Moravcsik, E. A. 2006. *An Introduction to Syntactic Theory*. London: Continuum.
- Moravcsik, E. To Appear. Part-whole relations in syntax. In *The Handbook of Mereology*. H. Burkhardt, J. Seibt & G. Imaguire (eds).
- Nichols, J. 1986. Head-marking and dependent-marking languages. *Language* 62 (1): 56–119.
- Postal, P. 1974. *On Raising. One Rule of English and Its Theoretical Implications*. Cambridge, MA: MIT Press.
- Radford, A. 1981. *Transformational Syntax. A Student's Guide to Chomsky's Extended Standard Theory*. Cambridge: Cambridge University Press.
- Rákosi, C. 2005. *Nyelvészet és Argumentációelmélet*. [Linguistics and Argumentation Theory]. Debrecen: DE Kossuth Egyetemi Kiadó.
- Sadock, J. M. 1991. *Autolexical Grammar. A Theory of Parallel Grammatical Representations*. Chicago: University of Chicago Press.
- Sharpe, D., Eakin, L., Saragovi, C. & Macnamara, J. 1996. Resolving apparent contradictions: Adults' and preschoolers' ability to cope with non-classical negation. *Journal of Child Language* 23: 675–691.
- Simon, H. A. 1996. *The Sciences of the Artificial*. Cambridge, MA: MIT Press.
- Speas, M. (ed.) 1990. *Phrase Structure in Natural Language*. Dordrecht: Kluwer.
- Tsohatzidis, S. L. (ed.) 1990. *Meanings and Prototypes*. London: Routledge.
- Tversky, B. 1990. Where partonomies and taxonomies meet. In *Meanings and Prototypes*, S. L. Tsohatzidis (ed.), 334–344. London: Routledge.
- Uexhuell, J. von. 1921. *Umwelt und Innenleben der Tiere*. Berlin: Springer.
- Varzi, A. C. No Date. Mereology. In *Stanford Encyclopedia of Philosophy*. (internet).

- Vennemann, T. & Nierfeld, G. 1973. Explanation in syntax. In *Syntax and Semantics*, Volume 2, J. P. Kimball (ed.), 1–50. New York: Seminar Press.
- Wertheimer, M. 1938. Laws of organization in perceptual forms. In *A Source Book of Gestalt Psychology*, W. D. Ellis (ed.), 71–87. New York: Harcourt, Brace and Company.
- Winston, M. E., Chaffin, R. & Herrmann, D. 1987. A taxonomy of part-whole relations. *Cognitive Science* 11: 417–444.

PART IV

Language, embodiment and cognition

Theory and application

Language as a biocultural niche and social institution

Chris Sinha

Grammars ... refer to real structures, though not to psychologically real structures in the processing sense ... a grammar is a description of our *knowledge of a social institution* – the language – and because of this basis in social or institutional reality, rather than in cognitive functioning, grammars and psychological processes have no more than the loose relationships they appear, in fact, to have. The role of grammar during speech programming is analogous to the role of other social institutions during individual behavior. This role is to define and evaluate the behavior of individuals. It is not to cause the behavior.

(McNeill 1979: 293)

1. Introduction: Language, culture and nature

The place of language in nature and culture is one of the abiding problems of all the language sciences, of which linguistics is but one. Language is at the heart of what it means to be human – indeed it has long been held that language is both *essential* to our humanity and *unique* to our species. Descartes famously argued that language is essentially human *because* it is an expression of uniquely human, universal reason. It is reason, he maintained, that distinguishes humans from animals, which are no more than merely complex machines:

For we can easily understand a machine's being constituted so that it can utter words, and even emit some responses to action on it of a corporeal kind, which brings about a change in its organs ... But it never happens that it arranges its speech in various ways, in order to reply appropriately to everything that may be said in its presence ... For while reason is a universal instrument that can serve for all contingencies, these organs have need for some special adaptation for every particular action. (Descartes 1911 [1637]: 116)

The Cartesian position was the precursor of contemporary theories of the universality and innateness of the human language faculty, although Descartes would perhaps have questioned the conceptualization of this faculty in terms of modular “mental organs” (Chomsky 1968; Pinker 1994).

Later Enlightenment philosophers maintained, on the contrary, that language, as an attribute of social association, culture and civilization, was in large part *responsible* for

human reason and for what came to be called the “higher mental processes.” Condillac, for example, wrote that:

The resemblance between animals and ourselves proves that they have minds, and the difference between them and us proves that their minds are inferior to ours ... the mental operations of animals are limited to perception, consciousness, attention, reminiscence and imagination not under their control, while ours include other operations whose origin I am about to make clear ... If contemplation consists in preserving perceptions, then before the use of institutional signs, it is merely outside our control; but if it consists in preserving the signs themselves, it has no function at all. So long as imagination, contemplation and memory are unused, or as long as imagination and contemplation operate outside our control, we cannot direct our attention as we please ... But when someone begins attaching ideas to signs of his own choosing, we see his memory begin to form ... Later, he acquires much greater control over his imagination as he invents more signs, for he has many more ways of using it. (Condillac 1987 [1746]:459)

Condillac rejected the nativism of Descartes, and his ideas not only draw upon the empiricism of Locke and Hume, but also anticipate Vygotsky’s notion that human higher mental processes are dependent upon their semiotic, and especially linguistic, mediation (Vygotsky 1978). Language, in his view, *makes* us human. As Condillac’s follower Itard – educator of “Victor”, the Wild Boy of Aveyron – wrote, “man is inferior to a large number of animals in the pure state of nature ... the moral superiority said to be natural to man is only the result of civilization, which raises him above other animals by a great and powerful force.” (Cited in Lane 1977: 129).

Descartes and Condillac agreed, then, that language distinguishes humans from other creatures; their disagreement was over whether it *expresses* or *enables* this difference, whether language is to be seen as primarily an expression of an innate faculty of mind, or as primarily a vehicle of social life and social interaction that enables and constitutes the uniqueness of the human mind. In this classical debate, one which defined not only the Enlightenment but Western thought about language down to the present day, we see the alignment of a number of dichotomous categories: human *vs* non-human; rational *vs* non-rational; culture *vs* nature.

Language, from this point of view, can belong *either* (following Condillac) to culture, making possible the “higher” faculties of human beings, *or* (following Decartes) to (higher, human) nature – defined in opposition to “lower” animal nature. Either way, the uniqueness of language is assumed to mirror the uniqueness of humans, interpreted in terms either of our unique nature, or, as we now would say, genetic make-up; or of the uniqueness of human culture. With the waning in recent years of the strong nativist program, it has become more usual to accept that the acquisition of language in children, like other aspects of development, is based upon epigenetic interactions between what is innate and what is available in the environment (Sinha 1988, 2004). In what sense might it also be possible to recast our ideas about the very *ontology* of language, in such a way that it comes to be seen as *both* a human socio-cultural form *and* intrinsic to human biology? And how, in its turn, might such a new, synthetic “biocultural” view of language affect our ideas about language processing and language learning?

2. Evolution, ecological niches and animal artefacts

A new synthesis is necessary because the assumptions guiding the thinking of both Cartesian, nativist and Condillacian, culturalist accounts are proving, in the light of 21st century research, to be fundamentally flawed. For the neo-Cartesians, human distinctiveness is to be sought in the genes, from which stem all human attributes, including cultural forms; for the neo-Condillacians, in the existence of human culture, a unique human attribute that, according to some, has rendered the very concept of a “human nature” invalid.¹ Both positions are predicated on an assumption of human uniqueness, of either genes or culture, and this assumption has motivated the traditional paradigm of language and its learning.

In the traditional paradigm, as we have seen, language is seen as part of *either* unique human nature (nativism), *or* unique human culture (environmentalism), and language learning is viewed as the exposure of the learner to an external “input” to be internalized. It is this model that, I argue, must be superseded, as a result of recent advances in biological sciences, advances that confront the traditional paradigm in the human sciences with a striking and challenging paradox.

The paradox is one of discontinuity in continuity. On the one hand, the biological characteristics of the human species display no dramatic discontinuities with those of other species; yet, on the other, human cognitive capacities, and human cultural constructions, appear from our current vantage point to be as exceptional in the living world as they did to Descartes. It can, of course, be argued that the cognitive and cultural discontinuity is merely a symptom of a gap in the available evidence – there are, after all, no living representatives of the human lineage since it diverged from the ancestors of our closest living primate relatives. If there were, the discontinuity would, perhaps, prove to be an illusion. Even so, it is hard to resist the conviction that, however extended the event, or sequence of events in evolutionary time, “something happened” involving language that radically transformed the evolving mind, and this transformation poses a profound and complex problem for both biological and social theory.

To begin with continuity: Darwin’s refutation of the idea that the human species is *essentially* different, in biological constitution and evolutionary history, from other species received, in the closing years of the last century, strong confirmation in two very different domains. Succinctly stated, neither genes nor culture, singly, can account for what, if anything, makes humans different from other species.

There is no evidence of dramatic genetic discontinuity between humans and their closest primate relatives, chimpanzees. The two species share, even on the most conservative estimate, about 95% of their genetic material (Britten 2002). Taken together with initial results of the human genome project, this suggests that whatever cognitive capacities distinguish the human species from other closely related species are unlikely to be attributable to dedicated genetic material available for directly coding such capacities. This

1. Malson (1972:9) (in his Introduction to Itard’s text) pursued the environmentalist direction of Condillac to its logical, if extreme conclusion, writing that “The idea that man has no nature is now beyond dispute. He has or rather is a history.”

does not mean that there is no genetic component of specifically human capacities. It does mean that the ascription of differences between the cognitive capacities of humans and those of non-humans to interspecies genetic differences *alone* is likely to be false. This is bad news for nativist modularity theories.

The news for those who would argue that what is unique about humans is the capacity for culture, a favoured hypothesis for generations past of social anthropologists, is hardly better. Culture can minimally be defined as the existence of intra-species group differences in behavioural patterns and repertoires, which are not directly determined by ecological circumstances (such as the availability of particular resources employed in the differing behavioural repertoires), and which are learned and transmitted across generations. On this definition, there is ample evidence of cultural differences in foraging strategies, tool use, and social behaviours in chimpanzees (Whiten et al. 1999; de Waal 2001). Such a definition will also qualify, for example, epigenetically learned intra-species dialect differences between songbird communities as cultural and culturally transmitted behaviour (Marler and Peters 1982). Again, this does not mean that there is *no* cultural foundation for uniquely human cognitive capacities; rather, it suggests that human culture, from an evolutionary and developmental point of view, must be treated as *explicandum* as much as *explicans*.

What is needed, it seems, is a theoretical apparatus capable of integrating culture and biology. One version of such integration, in which culture is analyzed as quasi-heritable units (“memes” or “culturgens”) has been argued for by sociobiologists such as Richard Dawkins (1976) and Edward O. Wilson (1998). Such accounts, however, have often been criticized for their reductionism, and recent biological theory suggests that the relation between biology and culture is far more of a two-way street than was ever envisaged by sociobiology. Far from eliminating culture by absorbing it into the genotype, some biologists are increasingly acknowledging the role of culture in shaping the evolutionary process *at the genetic level*, by the construction of new selective environments. Current developments in theoretical biology, amongst which the “niche construction theory” of Laland et al. (2000) is particularly significant, extend and modify the Neo-Darwinian synthesis that dominated 20th century biology by incorporating an ecological dimension that, I shall argue, proves to be particularly important for understanding human cognitive and linguistic evolution.

First, however, I outline (in a simplified fashion) the premises of, and the outstanding problems with, the Neo-Darwinian synthesis unifying Darwin’s theory of natural and sexual selection with modern population genetics.² In the Neo-Darwinian synthesis, the *unit of selection* (*what* is selected) is the gene, or more specifically alternative variants (alleles) of the “same” genes. The *agent of selection* (*what does* the selecting) is the extra-organismic environment, including (a) the inanimate surround, (b) other species (a and b together being the basis of natural selection), and (c) (subpopulations of) genes of the same species (the basis of sexual and kin selection). The relevant *attribute* upon which

2. As will become clear, there is no question here of challenging the overall Darwinian framework of evolutionary theory, but rather of questioning the premises and methodological stances of what is usually referred to as the Neo-Darwinian “modern synthesis”.

selection works (what is selected *for*) is any genetically transmitted trait. The mechanism of selection determines the differential reproductive success of the gene (allele) within the population of interacting genes, and thus the frequency distributions of genes and traits in the population. This model, when appropriately formalized, can be extended, as we shall see, by including cultural traits in the environment, that act as “amplifiers” on the selection of genetic variation: this is known as the theory of gene-culture coevolution (Lumsden and Wilson 1981).

The core issues at the heart of the problems besetting the Neo-Darwinian synthesis can be briefly summarized. First, genes do not come singly, but as combinations (genotypes), packaged in organisms (phenotypes). It is this distinction that Dawkins (1976) recasts as a distinction between the “replicator” (that which is copied), and the “vehicle” (that which embodies the genotypic collection of replicators, and interacts with the environment). For Dawkins, it is only the gene that is actually copied, and therefore he identifies the gene (unit of selection) as the replicator, and the phenotype as a mere “vehicle” for the replicator.

However, it is *organisms*, not genes, that are subject to *direct* selection pressures in terms of those traits conferring fitness. The organism level of biological organization receives scant attention in population genetics but, even granted that the gene is the *unit* of selection, it is the organism that must be considered as the *site* of selection. Organisms, in most (though not all) cases, can be regarded as morphological individuals. However, the actual *process* of selection by an “agent” occurs in relation to the *functioning, behaving* organism. It was for this reason that Jean Piaget upheld the leading role of behaviour in evolution (Piaget 1979). In the light of this, it may be (and frequently has been) questioned to what extent it remains legitimate to identify the “replicator” with the genetic unit of selection. Even if the DNA-based biochemical replicator is the gene, the *evolutionary dynamic* of replication-plus-selection should, it can be argued, more profitably be identified with the entire complex of the *site of selection*, which is the *active organism in its ecological niche*.

Ecologists emphasize that species shape, as well as being shaped by, their niches. Organismic behaviours may eventuate in significant *transformations* of the very environment to which the organism must adapt. A simple example (from Sinha 1988: 136) is the following: “A ‘path’ may ... be an unintended consequence of locomotion from one place to another, but it is, nevertheless, a useful one ... such shaping ... can [however] introduce distal consequences – food shortage, erosion, pollution, competition with other species – which are outside the initial circuit of adaptation” (see also Costall 2004). In many cases, however, a process of positive feedback will occur in which organism and environment are in a complementary relationship, each shaping the other. An oft-cited example is the hoof of the horse, and its adaptation to the grassland steppe whose ecological characteristics the horse, through its own motion through the landscape, reproduces.

In a subset of such cases, the resulting niche can be seen not merely as a contingent consequence of behaviour, but as an *animal artefact*, inasmuch as phenotypic individuals are genetically, morphologically and behaviourally *adapted* to the *production* of specific niches which are integral to the survival and/or reproduction strategy of the species. Examples of such artefactual niches are the nests of bower birds, and the dams of beavers.

The male bower bird builds and decorates an elaborate nest (bower) to attract females, using attractive objects such as flowers, shells and leaves. The bower forms an integral part of the male's mating display, and sexual selection by the female is based upon the bower as much as upon the behavioural display of the male. Beavers construct, through coordinated and collaborative behaviour, dams that serve both as a defence against predators, and as a means to enhance the availability of food. The dams of beavers not only serve as a constructed, artefactual niche for beavers themselves, but also reproduce the wetland ecology in which many other species thrive. As a final example of the significance of animal artefacts, we can mention the termite mound, whose material structure is not only integral to the reproductive strategy of this species of social insect, but also constitutes the morphological structure of the colony as a "group organism".

In each of these cases, the behavioural repertoire of the species includes behaviours that are specifically adapted to the making of artefactual niches, and these in turn support other behavioural strategies. The artefactual niche in many cases can be regarded as an extension of either a behavioural repertoire (e.g. male mating display) or of the organism's morphology (e.g. the bower bird's bower as functionally equivalent, as an indicator of fitness, to the tail of the peacock). Indeed, we can ask if it might be fruitful to consider certain species-specific behavioural repertoires, such as birdsong, to be kinds of animal artefacts, inasmuch the song of the adults provides a niche within which the singing behaviour is learned (see below). It can be argued that the designation of "artefact" should be reserved for more or less enduring, constructed material structures. Even if we accept this, it can still be argued that specialized behavioural repertoires constitute *biocultural niches* which are functionally analogous to animal artefacts. If so, human natural languages can also be viewed as species-specific biocultural niches.

3. Culture as constructed affordances and the human semiosphere

In the ecological psychology of James J. Gibson (Gibson 1979), a key role is played by *affordances*, properties of the ecological niche affording or supporting specific kinds of action made possible by the motor system and morphology of the animal. Such actions are both species-typical (though not necessarily species unique) and adaptive. Because affordances, Gibson maintained, are directly perceived, the phenomenal world of the animal is meaningful, in that it potentiates the activation of perception-action circuits: objects present themselves as edible, climb-able, graspable and so forth.

Gibson neglected, however, to note the crucial importance of the fact that some affordances are *constructed* by the animal itself. Artefactual niches are adaptive precisely because of the behaviours and strategies that they afford – nests are for nesting, and burrows are for burrowing. In such cases, the *site of selection* is no longer just the organism, but the organism in its self-constructed niche: the organism/niche coupling or *organism plus artefact*.

What are the implications of this for the Neo-Darwinian synthesis? A conservative reading would be that the only modification required is that the phenotype, or "vehicle", be extended to incorporate the artefactual niche. This is, indeed, the interpretation fa-

voured by Dawkins (1982), who employs the terminology of the “extended phenotype.” Under this interpretation, the “replicator” remains the gene, and only the gene. However, it is not only the gene that is copied or replicated. In fact, the artefactual niche too is both reproduced across generations, and serves as a fundamental precondition for genetic replication. The artefactual niche is thus *both a consequence of and an agent in natural and/or sexual selection*, and must then be seen as a key ingredient of the evolution of the species-typical genotype.

It seems, therefore, that the integration of ecological considerations into evolutionary theory, and specifically the existence of animal artefactual niches, further undermines the hard and fast distinction between germ-line and soma, genotype and phenotype, “replicator” and “vehicle”. In fact it makes better sense to say that, even granted that the *unit* of Darwinian selection remains the gene (allele), the “replicator” includes both the artefactual niche, and the niche-adaptive behavioural repertoire of the animal. Such considerations lead us back to Piaget’s more general proposition that behaviour is the leading edge and motor of evolution, prompting the conclusion (anticipated above) that the identification by Dawkins of the “replicator” with the *unit* of selection (the gene, or its hypothesized cultural analogue, the “meme”) is deeply flawed, and that replication can as well or better be considered as a property of the entire *site* of selection.

At this point, it is useful to make a brief critical detour to re-examine Neo-Darwinist theories of gene-culture co-evolution such as that of Lumsden and Wilson (1981). Such accounts presuppose a functional parallelism between units of biological replication and units of cultural replication (“memes” or “culturgens”); and treat the latter as being structured as human behavioural variants analogous to gene alleles. From this perspective, ethnographic variation is analyzable in terms of aggregate properties of human populations. The Lumsden-Wilson theory has been criticized for making “the reductionist assumption that the characteristics of a society can be understood as simply the sum of the characteristics of the individuals of that society” (Alper and Lange 1981: 3976), and for having no place for *emergent properties* of societies. To lend further force to this critical evaluation, I attempt in a subsequent section to formally specify, for human cultures, such emergent properties in terms of a semiotically defined ontology of the social. For the time being, we can simply note that the Lumsden-Wilson theory *presupposes* an ontological distinction between gene and meme, nature and culture, without either explaining this distinction, or theoretically motivating the functional parallelism that is proposed to exist between the units of selection in the domains of biology and culture. In summary, the reductionist inadequacy of Neo-Darwinist theories of gene-culture coevolution consists in their recapitulating the failure of Neo-Darwinism to adequately treat the emergent properties of organisms, in their failure to adequately treat the emergent properties of socio-cultural formations.

The critical considerations outlined above have led to a more radical formulation of coevolution, advanced by Laland, Odling-Smee and Feldman (2000).³ A particular role is played in Laland et al.’s (2000: 144) theory by genotype/niche combinations labeled by

3. Here it is important again to emphasize that the theory advanced by Laland et al., while radically departing from Neo-Darwinism, remains Darwinian in the wider sense.

“phenogenotypes.” A phenogenotype can be defined as a class of organisms in a bound (though not necessarily genetically determined) relationship with some aspect of a self-constructed environmental niche.

Laland et al. (2000:132) criticize the “human-centred” perspective of previous accounts of gene-culture coevolution, emphasizing that many *non*-human species behaviourally co-direct genetic evolution through niche construction. This point is important, because it situates the role of culture in human evolution within the wider class of processes, outlined in the previous section, involving adaptation to artefactual niches such as nests, dams, mounds, and burrows. Laland et al.’s model, then, is a general one, not confined to human culture and evolution. They acknowledge, however, that humans are “unique in their extraordinary capacity for culture” (ibid.: 133). I interpret this to mean primarily that human cultures are unique in some fundamental respect, that is they are different (perhaps discontinuously) from the cultures of other species; and secondarily that the capacity for creating, acquiring, and transmitting cultural forms is uniquely developed (though clearly not unique) in humans.

One evident discontinuity between human and non-human cultures is that human cultures are linguistic; and the capacity for human cultural acquisition and transmission is mediated by the unique human language capacity. The nativist modularist account of this capacity proposes its inscription in the human genotype, a hypothesis vulnerable to many objections, including the difficulty stated above of locating this profound discontinuity in the continuous landscape of the primate genome. An alternative account, along the lines of the co-evolutionary theory of Laland et al. (2000), would view the human language capacity as phenogenotypic. Language, in this account, is an artefactual niche, and the capacity to acquire and use it involves the evolution and replication of a phenogenotypic “biocultural complex” (Laland et al. 2000: 144).

Such an account does not require the organism to possess an internal model of the grammar of a language to account for language acquisition, any more than the building of a nest requires a prior internal model of the nest. The grammar of the language is *in the language*, just as the structure of the nest is in the nest. The capacity for language is thus a cognitive-behavioural relationship between language user and the constituents of language, just as the capacity for building a nest is a cognitive-behavioral relationship between the builder and the constituents of the nest; and it is this *relationship* that, in each case, has been selected for in evolution. This account is thus compatible with usage-based, cognitive functional theories of language and language acquisition (Tomasello 1998, 2003).

The artefactual niche of language is culturally situated, that is, it is dynamically embedded within the entire semiotic biocultural complex that includes other symbolic and non-symbolic artefacts. This biocultural complex we can, to use the terminology of the Russian semiotician Yuri Lotman (e.g. Lotman 1990), designate as the human *semiosphere*, the constructed, meaningful environment that is reproduced down the human generations along with the human organism itself. It is crucial to appreciate, in this context, that the semiosphere, like other animal artefactual niches, is not merely a constituent of *what* is reproduced, but is also the fundamental mechanism in the *process* of reproduction and transmission.

Because of its pre-eminence in mediating both cultural reproduction and individual cognitive processes, language is the primary and most distinctive constituent of the human semiosphere. The class of organisms with the language capacity (normally developing humans) can thus be theorized as a phenogenotypic replicator, systemically associated with a wider biocultural complex of symbolic and constructive cognitive capacities, also of a phenogenotypic nature; and individual language acquisition and use is situated in the contexts of actuation of these inter-related capacities. This account accords with the view that what makes humans unique is not an innate language acquisition device plus a variety of other species-specific innate cognitive modules, but a generalized semiotic or symbolic capacity (Piaget 1945; Deacon 1998; Zlatev et al. 2006); epigenetically developed from a suite of cognitive capacities largely shared with other species, but attaining higher levels of organization in humans.

4. The evolution of complexity: Emergence and epigenesis

It was noted above that one of the criticisms made by Alper and Lange (1981) of the Lumsden-Wilson theory of gene-culture coevolution was its inability to account for emergent properties of human society. The term *emergence* is commonly used to mean the evolution and development of new properties and/or levels of organization of behavioural and cognitive systems as a consequence of the operation or cooperation of simpler processes. Emergence is a consequence of the evolutionary and developmental process of elaboration, involving an increase in the complexity of organism, behaviour and cognition. The adaptation of organismic behaviour to constructed, artefactual niches offers clear examples of phylogenetic elaboration leading to emergence. For example, the female bower bird's behavioural repertoire for evaluating the fitness of a prospective mate has extended (elaborated) its scope from evaluating the ritualistic behaviour of the male suitor to evaluating the bower that he constructs, entraining a more complex organismic "level" with emergent properties (the bower as a constituent of the phenogenotypic replicator).

Emergence as an evolutionary process can be conceptualized as "locking" elaboration in new adaptive circuits, in a way similar to the "ratchet effect" discussed for human cultural evolution by Tomasello (1999). Indeed, from the biocultural perspective, Tomasello's ratchet effect is but one, albeit exceptionally dramatic, case of a wider, not specifically human, phenomenon of emergence. If it is elaboration that lends directionality (through complexification) to "local" evolutionary processes, it can be hypothesized that it is emergence that underlies the global trend of evolution towards greater complexity. Although it is correct to reject teleological *explanations* for Darwinian evolution, a kind of teleology of process (as Piaget recognized) is a consequence of the locking and reproduction of elaboration through emergence.

Emergence is also characteristic of ontogenetic development, including cognitive development, and has been advanced as an alternative to nativist accounts of language development (MacWhinney 1999). The developmental biological (and psychobiological) mechanism underlying ontogenetic emergence is *epigenesis*.

Contemporary theories of epigenesis in biological and psychological development build upon the pioneering accounts of Waddington (1975) and Piaget (1979). Epigenetic naturalism (Sinha 1988) proposes a constructivist account of the interaction between the genotype and its somatic and extra-somatic environment in organismic development. The claim that such an interaction exists is, as such, trivial and undisputed, since everyone agrees that phenotype is co-determined by genes and environment. There are, however, two particularly important characteristics of the theory of epigenesis that I wish to highlight.

The first is that the role of the environmental factors is *constructive* in addition to being *selective*. Nativist approaches to the developmental interaction between genotype and environment stress the role of specific input either in permitting a developmental process to unfold, or in parametrically selecting a particular variant of development. An example of the former would be phenomena such as “imprinting”, where an innate and fully endogenous process of development is “triggered” by an environmental event during a critical developmental window. An example of the latter would be the role hypothesized by generative linguists to be played by typological characteristics of target languages in setting parameters and thereby permitting the child non-inductively to acquire the grammar of the target language (Chomsky 2000). In neither of these cases does the environmental information add any emergent *level of organization* to the genetically coded information. That is to say, the alternate pathways along which the behaviour develops, and its terminal structural complexity, are assumed already to be directly encoded in genes.⁴

By contrast, in epigenesis the developmental pathway and final structure of the behaviour that develops are a consequence as much of the environmental information as of the genetically encoded information. For example, the development of birdsong seems to involve reproduction by imitative epigenetic learning, rather than selection from amongst pre-established alternatives (Marler and Peters 1982). Fledglings not exposed to a model do develop birdsong, but it is impoverished or unelaborated relative to that of those individuals developing in a normal environment in which models are available.

4. It is noteworthy that the Lumsden-Wilson account of gene-culture coevolution, though it employs a terminology of epigenesis, does so in a way that is more reminiscent of the Chomskyan notion of parameterization. The Lumsden-Wilson theory envisages two (or, in principle, more) alternate traits distributed in a population, with the possibility of epigenetic “switching” between these traits, which could therefore be alternate expressions of the same gene(s). They hypothesize that an interaction of “epigenetic rule”-encoding genes, and environment (in particular, existing trait frequencies in the population), is responsible for the development in any given individual of one or other trait. Alper and Lange (1981: 3976), whose critical assessment of the Lumsden-Wilson theory I have already cited, claim that there is “absolutely no evidence that any genes of this type exist.” This author is not competent to pronounce on the biological facts of this disputed point. However, it should be noted that (a) Lumsden and Wilson’s “code-switching probability” genes are substantially different from the regulatory genes discussed below, and (b) their model continues to fix the stochastic space of possible developmental outcomes in the genes, rather than this itself being co-determined by genes and environment. Given these considerations, it could be said that the Lumsden-Wilson gene-culture coevolution model (intended as a substantial revision of standard sociobiological theory) severely circumscribes the role of epigenesis in such a way as to call into question whether it is really epigenetic at all.

The second key characteristic of epigenesis is, accordingly, that a genetically specified developmental envelope or window specifies an initial behavioural (or perceptual) repertoire that is subsequently *elaborated* through experience of a relevant environment. This process of elaboration is directional (see below), and once it has taken place the initial plasticity of the embryonic, or unelaborated, repertoire is lost. A typical example is the development in human infancy of speech sound perception, in which the “universal” initial processor is transformed into a “language-specific” processor in a process that is probably analogous with that of the development of birdsong. We can note here that an epigenetic account of this process differs from a nativist, parameter-setting process inasmuch as no assumption is made that the infant brain is innately equipped with an inventory of all possible natural language phonemes. Equally, however, it differs from a classical learning account, inasmuch as epigenesis depends upon the elaboration of an initial repertoire which itself is not learned, in a process which cannot be re-run – the initial, unelaborated capacity cannot be re-accessed after the epigenetic developmental process has taken place, as all second language learners come rapidly to realise. In other words, the process of developmental elaboration implies in epigenetic development a transition from relative plasticity and informational openness to relative rigidity and informational closure.

Epigenesis is a developmental process whose genotypic distribution can itself be selected, through the standard Darwinian mechanisms of natural and sexual selection. The onset and closure of periods of plasticity and informational openness is under the control of regulatory genes, which are responsible for the timing and sequencing of all developmental processes from embryonic to mature organismic stages. Although all mammals display epigenetic features in various domains of development, especially during embryogenesis, the “weighting” of behavioural and cognitive development towards epigenetic processes seems to be markedly greater in more complex organisms than less complex ones. As an example, we can cite the well-known difference between the capacities of human-enculturated vs non-enculturated apes for symbol learning. The enhanced learning capacity of the enculturated apes must be due to their developmental environment, but the developmental environment would not make a difference if its relevant features were not available for assimilation into epigenetic construction processes (and indeed in lower non-human primate species no such differences have been observed). Epigenesis, then, seems to be a key mechanism in enabling individual organisms to acquire and exploit emergent complexity in phenogenotypic couplings. It is epigenesis, and in particular selectively *augmented* epigenesis, that serves as the crucial bridge between parallel processes of emergence in phylogeny and ontogeny. Epigenesis, I would suggest, is the key stabilizing process through which cultural and cognitive elaboration are emergently, dynamically and concertedly locked.⁵

5. This is something of an oversimplification, since the stabilizing role accorded here to epigenesis also involves the canalization of phylogenesis through “Baldwin effects” and genetic assimilation (Sinha 1988: 137–138). However, this detail is not crucial to the account presented here, and the precise status and nature of “Baldwin effects” is a strongly contested issue (Deacon 2004).

Epigenesis may be analyzed at different levels, from the cellular to the cognitive. The neurobiological implementation of epigenetic processes at the cognitive and behavioural level is “Neural Darwinism,” the selective stabilization of waves of periodically proliferating synaptic connections during ontogenesis (Changeux 1985), which is, as we might expect, especially marked in the developing human organism, extending through adolescence and young adulthood. Why should humans be, more than any other, the “epigenetic species,” and how is this related to humans being the “symbolic species” (Deacon 1998)?

The answer to this question, I suggest, is that augmented epigenesis is advantageous for organisms in which phenogenotypic organism-niche couplings are both frequent and variable, which is a good enough general description of the human cultural organism. Regulatory genes augmenting epigenetic openness can therefore be expected to have been phenogenotypically selected for in the human genome, permitting further adaptive selection for domain-specific learning in the semiotic biocultural complex, in particular for language. Note, however, that in an epigenetic perspective, any developmental predisposition for learning language is unlikely either to involve direct coding of, or to be dedicated exclusively to, linguistic structure (Mueller 1996).

The account I have offered revolves around the proposition that the evolutionary elaboration and epigenetic stabilization of the phenogenotypic semiosphere introduced the discontinuity characterizing both human culture and human cognition. Signs are both transformative cognitive tools, and constitutive of specifically human cultural ecologies. The semiotic capacity is hypothesized to have triggered transformative effects across all or most cognitive domains, thereby potentiating human symbolic cultures, which constitute the biocultural niche complexes in which human cultural innovation and transmission occur. The semiotic capacity is the explanatory link binding what is unique to human cognition with what is unique to human culture, bridging the biological with the social and human sciences in the evolutionary and developmental science of human cognition and language. It is to the social and semiotic ontology of language that I turn in the next section.

5. Language as a social fact and social institution

I begin by summarizing two theories of social ontology, classical and modern, separated in time by a century. These are the theories of, respectively, the sociologist Emile Durkheim (1895) and the philosopher John Searle (1995). Probably not by coincidence, the theories employ the same terminology of social facts and social institutions, although Searle nowhere cites Durkheim.

Durkheim, a founding father of social theory, attempted a theoretical and methodological clarification of social science and its object. This object he stipulated to be the domain of *social facts*, which he described as “a category of facts which present very special characteristics: they consist of manners of acting, thinking, and feeling external to the individual, which are invested with a coercive power by virtue of which they exercise control over him” (Durkheim 1982 [1895]).

Social facts, for Durkheim, are not merely aggregates of the individual cognitive representations of them by the subjects that are regulated or “coerced” by the social facts, since for each individual subject the social fact presents itself as a part of an out-there, already given *objective* reality. The objectivity of social facts consists, for Durkheim, in the fact they are independent of any single individual’s thoughts or will. As Jones (1986: 61) puts it, “it is precisely this property of resistance to the action of individual wills which characterizes social facts. The most basic rule of all sociological method, Durkheim thus concluded, is to *treat social facts as things*.” Durkheim’s treatment of social facts consists therefore in, first, an ontological proposition, that social facts are irreducible to biological or psychological facts (or structures or processes); coupled with, second, an epistemological and methodological proposition regarding their treatment: as *objects* of a particular kind, whose determinate nature consists in their “coercion” of conduct.

Durkheim has often been criticized for the breadth and vagueness of his notion of “social fact”. A particularly problematic aspect of his theory is that, in counterposing “social facts” to “individual conscience” (or mind), he sometimes identified the former with “states of the collective conscience”. Some social psychologists (e.g. Moscovici 2000) have followed this direction in constructing a theory of “social representations”, but critics have claimed that Durkheim sympathized with a view of society as a kind of super-organic “collective personality”.

Whether Durkheim believed in a “collective mind” or not, such a notion is not only scientifically untenable, it is unnecessary. I propose that a social fact can most simply be defined as something regulating an aspect of conduct which requires the *participation* (Goodwin and Goodwin 2004) of more than one individual. This “something” may be a codified law, a norm, an institution, a rule in the Wittgensteinian sense, or a canon of interpretation. A natural language, therefore, qualifies as a social fact (or institution) under this reading of Durkheim’s theory. Social facts, for mature human beings, are objects of *common knowledge*; language is a prime example of this (Lewis 1969; Ikonen 1983; Clark 1996). However, the social fact itself is not the sum, average or common denominator of all the individual beliefs of participants (since it is, indeed, the *object* of these beliefs). Social facts, in this sense, are in some way prior to individual cognitions about them. Yet it cannot be claimed that social facts are *independent* of cognitions, in this case of social cognitions, since their normative status is dependent upon agreement in cognition. We shall return to this paradoxical problem in discussing Searle’s theory of social facts.

Social facts, for Durkheim, are constitutive of the domain of human social theory. Given that non-human species also display social behaviours, should we regard social facts as being uniquely human? Ethologists have pointed to the evolutionary roots of norms, rules and conventions in the ritualized displays that many species exhibit in, for example, mating and agonistic displays. Ritualization, in turn, can be regarded as falling under the definition of a biocultural niche as discussed above. If so, we could argue that social facts are biocultural niches regulating and sustaining, supporting and constraining, the participatory behaviour of more than one individual. This definition is entirely consonant with Durkheim’s view that social facts “consist of manners of acting, thinking, and feeling external to the individual, which are invested with a coercive power by virtue of which they exercise control over him.”

Under this interpretation, social facts would be seen as no more unique to humans than culture. Yet there is clearly something unique about human social facts. This uniqueness consists, surely, in the way in which social facts are cognitively constructed as *objects* of intersubjective common knowledge (and common emotional investment), so that they can be *known* in the way in which the rules of football, the laws of the land, or a family history may be known.

To bring some order into the definitions employed here, I will stipulate that the concept of “social fact” pertains to those biocultural niches which are of a fully normative nature; that is, those which not only regulate behaviour, but are known to do so, and knowledge of which (whether explicit or tacit) is essential to their regulative status. Social facts, on this definition, can only be *constructed* by human beings with a certain level of cognitive development, although the institutions that they construct may be participated in by animals which lack this cognitive status (e.g. prelinguistic infants in language practices, racehorses in horse races etc). Social facts, then, constitute an emergent ontological level within the wider category of biocultural niches, and one which is uniquely human.

Searle (1995) situates knowledge and belief at the heart of his account of social facts: “There are things that exist only because we believe them to exist. I am thinking of things like money, property, government, and marriages ... [such] Institutional facts are so called because they depend upon human institutions for their existence” (ibid.: 1–2). In an unfortunate inconsistency of terminology, Searle regards “institutional facts”, which seem to be more less equivalent to Durkheim’s social facts, as a subset of what he (Searle) calls “social facts”, which are basically all activities which involve participation in joint action: “I will henceforth use the expression ‘social fact’ to refer to any fact involving collective intentionality. So, for example, the fact that two people are going for a walk together is a social fact. A special subclass of social facts are institutional facts ... for example, the fact that this piece of paper is a twenty dollar bill is an institutional fact” (ibid.: 26).

Searle’s account of social or institutional facts (such as money) is that they depend upon collective agreement and knowledge that, under determinate rules, something *counts as* an instance of a social object. Hence, the general form of such rules is:

1. “X counts as Y in context C” (Searle 1995: 28).

Although he never uses the term, Searle’s definition is in effect a semiotic one, in that the “counting as” relationship is one of *meaning* or *signification*. The twenty dollar bill, for example, signifies a certain monetary value or equivalence. However, the relationship between the bill and its monetary value is not a fully-fledged *sign relationship*. The bill does not *represent* or *stand for* twenty dollars: it simply *is* twenty dollars, it is self-identical to its monetary exchange value. To clarify this difference, we can point out that the numeral 20 printed on the bill *stands for* (represents) the number twenty, but the bill itself does not represent, for example, twenty one dollar bills, but rather is *equivalent* to them in the value that it possesses, or counts as having.

Sinha (1988: 37) defines the pragmatic and semiotic conditions on *representation* as follows: “To represent something ... is to cause something else to stand for it, in such a way that *both* the relationship of ‘standing for’, *and* that which is intended to be represented, can be recognized.” (italics added). It must be emphasized that built into the conditions

on representation is a duality of cognition, paralleling the duality of sign structure (the conventional unity of signifying substance and its signification). Two cognitions are necessitated: the recognition of the sign relationship, and the recognition of what is signified. The “counting as” relationship, by contrast, has no such duality: to know that something counts as a particular object, however abstract or complex that object may be, it is necessary only to recognize it as a token of that category of objects.

What is necessary to grasp the “counting as” relationship is knowledge of the rules and norms that constitute the category (for example money, or a language). In one fundamental (if limited) sense, then, knowledge of a language is definitionally knowledge of what *counts as* a token of the language, and in order to know this, the knowing subject must necessarily know (in some way and to some degree) the rules of the language. It is this level of knowledge that is considered to be primary in generativist and other formalist theories of language, which attempt to elucidate the rules that constitute the full range of tokens for which it is the case that:

2. X counts as (a sentence) S in L (a language)

This definition does not, however, encompass in any way the *representational* function of language, its capacity to represent things (situations, events, actions, objects) *outside* of the formal context of L, that is, the world outside language. The knowledge constituting this *semantic* domain is governed, not by the “counting as” relationship and its conditions, but by the “standing for” relationship” and its conditions. This “standing for” relationship can be notated, in a way parallel with Searle’s notation of the “counting as” relationship, as follows:

3. S (a sign) stands for M (a message) in context C

However, the duality inherent in the conditions on representation (above) requires that this preliminary notation be expanded, to include knowledge on the part of the subject that S *counts as* a sign, or, more accurately, that a particular object counts as a signifier. This expansion yields:

4. [X counts as S and S stands for M] in C
Where X is a token of the class of signifiers in C

(4) is sufficiently general to cover all cases of sign use, including highly idiosyncratic and context bound cases, such as non-conventional gestures. We can now undertake a further expansion to specify cases in which a given sign is part of a sign *system*, shared by a particular community of users:

5. [X counts as S and S stands for M in C_s] for C_u
Where:
C_s = sign system
C_u = community of users

In the specific case of language, we can reduce the notion of a sign system shared by a community of users to the simple term L, language, thus:

6. $L = C_s$ for C_u

Now any grammatical and meaningful instance of language use can be annotated:

7. [X counts as S and S stands for M] in L

Note that, consistently with the approach of Cognitive Grammar (Langacker 1987), S (the signifier) is an expression at any level, sub-lexical, lexical or constructional; grammatical assemblies of signs are also signs.

The definition offered in (7) can thus be considered to be the notational reduction of the broader theoretical approach to language taken by cognitive and functional semantically based theories, and indeed by all linguistic theories that include representational meaning in the linguistic theory. It is clearly a more inclusive definition than the formal-sentential definition (2), reproduced here:

2. X counts as (a sentence) S in L

Defintion (7) is also, quite simply, more psychologically complete than (2): what we usually mean by “knowing a language” is the knowledge of *both* what counts as a token of the language, *and* what it means. In the concluding section, I will attempt to elucidate further just what is, and is not, necessary for such knowledge.

Before doing so, I pursue this formal-notational exercise further by exploring how the conjoint definitions of “counting as” and “standing for” can be employed to define the sub-systems of language as traditionally employed in linguistic theory.

Grammar (in the wide, cognitive grammar sense, including lexical form and phonology) can be defined as:

8. X counts as S in L

X is an instance of S, and S is a grammatical expression in L. The distinction between X and S is the distinction between, for example, phonetics and phonology.

Presupposing (8), **semantics** can be defined as:

9. S stands for M in L

This is the relation between, for example, word form and lexical entry or concept; or, more generally, between linguistic expression and linguist conceptualization.

Presupposing (9), **pragmatics** can be defined as:

10. S counts as A_s for Participants_(2 ... n) in C_d

Where:

A_s = Speech act (including reference)

C_d = Discourse context

Under this description, pragmatics is the closest of the linguistic subsystems to the “counting as” relationship. This accords with the intuition that pragmatics is not “systematic” in quite the same way as grammar and semantics; that speech acts are specifically linguistic instances of more general communicative acts (such as “threats” and “invitations” in both

human and non-human species); and that their interpretation is strongly dependent on gesture, prosody, posture, physical and linguistic context.

Having employed the notational formalism to distinguish the subsystems of language one from another, we can now re-assemble them to analyze the structure of particular utterances in their context.

11. [X counts as S and S stands for M] in L and S counts as A_s for Participants $(2 \dots n)$ in C_d

Such a re-assembly does not yet account for the *interaction* between semantics, pragmatics, extra-linguistic context and shared world knowledge in actual utterances. For example, if the utterance is “You really did well this time!”, and it is clear from the context that the speech act is one of ironic praise, the contextual meaning is “You did very badly”. Or, if the utterance is “The road meanders up the hill”, the contextual meaning is that the road has a winding path, not that the road is itself in motion (Talmy 1996). How can we capture such facts of language?

It seems impossible to do so without appealing to psychological processes such inference, default and prototypic reasoning, subjectivization and perspectivization. This is the fundamental insight which drives cognitive linguistics. If we wish to formalize this, it would look something like this:

12. [X counts as S and S stands for M] in L and S counts as A_s in $C_d \Rightarrow$ S counts as (having) M_c for H in C_d
 Where:
 M_c = Contextual meaning
 H = Hearer

This brings us back, in an intriguing hermeneutic circle, to Searle’s original definition of a social fact, and emphasizes the truism that, in the end, all meaning is contextual and situated. This does not, however, mean the same as saying that there are no institutionalized, relatively stable, relatively autonomous and systematic social facts; indeed, it is precisely this very relative stability and autonomy which constitutes the objectivity of social facts emphasized by Durkheim.

This objectivity is not to be confused with the *objectivism* of formal, truth conditional semantics. Amongst the advantages of the simple notational definitions developed here are:

1. The account of semantic meaning is underdetermined by this formulation. The semantic theory need not be truth-functional, but *is* (necessarily) conventional and normative (as indeed are all the subsystems).
2. Semantics is distinguished from pragmatics without necessitating a truth functional semantics.
3. Contextual dependence characterizes all subsystems, as well as the interactions between them, but does not erase the distinctions between them.

4. Language as a social institution has its own proper structure which necessitates, but is irreducible to, the intentionality of its users. Language, like all social facts, is an objectification of intersubjectivity, with an emergent structure relatively autonomous from the intentional states (such as mutual knowledge of the language) which are possessed by its users and "subjects". It is in this fact, and this fact alone, that the objectivity of language inheres.

6. Conclusion: Restricting the need to know

Biocultural niches are integral to the evolution of many species, including the human species. Radical nativist hypotheses of strong discontinuity between human and non-human genome are neither necessary nor plausible, if we view biocultural niches as constituents of phenotypic sites of Darwinian selection. The human semiosphere can be viewed as a species-specific biocultural niche, whose distinguishing feature consists in the elaboration and emergence of the semiotic function. This function, in turn, is constituted by the interplay and developmental interlacing of its two constituent semiotic relations, "counting as" and "standing for". It can be hypothesized that these two semiotic relations are evolutionary derivatives of, on the one hand, ritualization and, on the other, the evolution of symbols from signals (Sinha 2004). In both of them, the conventionalization of intersubjective participation in niche-regulated activities plays a central role (Sinha and Rodríguez 2008). Current evolutionary biological theory, including niche-dependent evolution and epigenesis, accounts for the continuity between human and non-human culture and cognition. The evolution of the human semiosphere, in which language as a biocultural niche is developmentally and processually interdependent with other artefactual supports for human social interaction and social practice (Sinha 2005), is what accounts for the discontinuity dividing human from non-human cognition and culture, and the evolutionary emergence of human social facts and social institutions. This discontinuity has been amplified by the consolidation, through language, of human culture as a fundamentally symbolic order.

From a biocultural perspective, the human language capacity, although it is almost certainly supported by genetic adaptations to maximize exploitation of the human biocultural niche, is not innate, but epigenetically developed. Language as a social institution comes to be known by language-acquiring infants, but the knowledge required is not that of a grammar as a formal object divorced from its semiotic function. Although there can be no scientific objection to the study of language as a purely formal system, insistence on the disciplinary autonomy and full explanatory adequacy of formal theories leads to a distorted picture of the human language capacity, and to unnecessarily constrained theories of language acquisition. If "knowledge of language" is restricted to knowledge of what counts as a grammatical sentence, not only is language itself as a semiotic system truncated and reduced, but the process of its acquisition is rendered incomprehensible. To fill this conceptual vacuum, innate knowledge of Universal Grammar is invoked.

The biocultural theory of language and its acquisition restores, quite literally, life to language, for far too long reduced to formal structures and operations. It suggests a picture of "knowledge of language" that is both richer, in one sense, and poorer, in another,

than that to which we have become accustomed from generative linguistics. It is richer because it incorporates meaning and context, the fundamental pillars supporting both language acquisition and language use. It is poorer because there is no longer a compelling reason to attribute a knowledge equivalent to the results of formal analysis to the learners and users of language. Simply stated, in the biocultural theory, *there is no mental grammar* isomorphic with autonomous grammar. Rather, grammar is *in language*, as a biocultural niche and social institution, just as the structure of the bower bird's nest is in the nest as a biocultural niche. The learner need not internalize a formal description of the structure in order to acquire the ability to *act* in it. Language is not an "input" to a processor or device, but a structured niche affording complex and semiotically mediated communication and cognition. Grammar is a social institution, normatively regulating linguistic practice, and it is the practical ability to adhere to its constraints and supports that is acquired by the language learner.

Acknowledgments

I wish to thank my colleagues Alan Costall and Jörg Zinken for their valuable comments on parts of this article. I am also grateful for the stimulating environment provided by my colleagues on SEDSU (Stages in the Evolution and Development of Sign Use), a collaborative project supported by the European Union 6th Framework programme "What it Means to be Human" (Zlatev et al. 2006).

References

- Alpert, J. & Lange, R. 1981. Lumsden-Wilson theory of gene-culture coevolution. *Proceedings of the National Academy of Sciences USA* 78: 3976–3979.
- Britten, R. 2002. Divergence between samples of chimpanzee and human DNA sequences is 5%, counting indels. *Proceedings of the National Academy of Sciences USA* 21: 13633–13635.
- Changeux, J.-P. 1985. *Neuronal Man: The Biology of Mind*. Oxford: Oxford University Press.
- Chomsky, N. 1968. *Language and Mind*. New York: Harcourt, Brace and World.
- Chomsky, N. 2000. *The Architecture of Language*. Oxford: Oxford University Press.
- Clark, H. H. 1996. *Using Language*. Cambridge: Cambridge University Press.
- Condillac, E. B., Abbé de 1987 [1746]. *Philosophical Writings*, Volume 2, transl. F. Philip. Hillsdale, NJ: Lawrence Erlbaum.
- Costall, A. 2004. From Darwin to Watson (and Cognitivism) and back again: the principle of animal-environment mutuality. *Behavior and Philosophy* 32: 179–195.
- Dawkins, R. 1976. *The Selfish Gene*. Oxford: Oxford University Press.
- Dawkins, R. 1982. *The Extended Phenotype*. Oxford: Oxford University Press.
- Deacon, T. 1998. *The Symbolic Species: The Co-Evolution of Language and the Human Brain*. Harmondsworth: Penguin.
- Deacon, T. 2004. Plenary lecture to the 1st International Conference on Language, Culture and Mind, University of Portsmouth, July.

- Descartes, R. 1911 [1637]. *Discourse on the Method of Rightly Conducting the Reason*. In *The Philosophical Works of Descartes*, E. S. Haldane & G. A. Ross (transl. and eds). Cambridge: Cambridge University Press.
- de Waal, F. 2001. *The Ape and the Sushi Master*. London: Allen Lane.
- Durkheim, E. 1895. *Les Règles de la méthode sociologique*. Paris: Alcan.
- Durkheim, E. 1982 [1895]. The Rules of Sociological Method. In *The Rules of Sociological Method and Selected Texts on Sociology and its Method*, S. Lukes (ed.), 29–163. London: Macmillan.
- Gibson, J. 1979. *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.
- Goodwin, C. & Goodwin, M. H. 2004. Participation. In *A Companion to Linguistic Anthropology*, A. Duranti (ed.), 222–244. Oxford: Blackwell.
- Itkonen, E. 1983. *Causality in Linguistic Theory*. London: Croom Helm.
- Jones, R. A. 1986. *Emile Durkheim: An Introduction to Four Major Works*. Beverley Hills, CA: Sage Publications.
- Laland, K. N., Odling-Smee, J. & Feldman, M. W. 2000. Niche construction, biological evolution, and cultural change. *Behavioral and Brain Sciences* 23: 131–175.
- Lane, H. 1977. *The Wild Boy of Aveyron*. London: Allen and Unwin.
- Langacker, R. W. 1987. *Foundations of Cognitive Grammar. Theoretical Prerequisites*. Volume 1. Stanford: Stanford University Press.
- Lewis, D. K. 1969. *Convention: A philosophical study*. Cambridge, MA: Harvard University Press.
- Lotman, Y. 1990. *Universe of the Mind: A semiotic theory of culture*. (Transl. Ann Shukman). New York: I.B. Tauris and Co. Ltd.
- Lumsden, C.J. & Wilson, E.O. 1981. *Genes, Mind and Culture: The Coevolutionary Process*. Cambridge, MA: Harvard University Press.
- MacWhinney, B. (ed.). 1999. *The Emergence of Language*. Mahwah, NJ: Lawrence Erlbaum.
- Malson, L. 1972. *Wolf Children and The Wild Boy of Aveyron*. London: NLB.
- Marler P. & Peters, S. 1982. Developmental overproduction and selective attrition: New processes in the epigenesis of birdsong. *Developmental Psychobiology* 15: 369–378.
- McNeill, D. 1979. *The Conceptual Basis of Language*. Hillsdale, NJ: Lawrence Erlbaum.
- Moscovici, S. 2000. *Social Representations*. Cambridge: Polity Press.
- Mueller, R.-A. 1996. Innateness, autonomy, universality? Neurobiological approaches to language. *Behavioral and Brain Sciences* 19: 611–675.
- Piaget J. 1945. *La Formation du symbole chez l'enfant*. Neuchatel: Delachaux and Niestlé.
- Piaget, J. 1979. *Behaviour and Evolution*. London: Routledge.
- Pinker, S. 1994. *The Language Instinct*. London: Penguin.
- Searle, J. 1995. *The Construction of Social Reality*. London: Allen Lane.
- Sinha C. 1988. *Language and Representation: A Socio-Naturalistic Approach to Human Development*. Hemel Hempstead: Harvester-Wheatstheaf.
- Sinha C. 2004. The evolution of language: From signals to symbols to system. In *Evolution of Communication Systems: A Comparative Approach*, D. K. Oller & U. Griebel (eds), 217–235. Cambridge, MA: MIT Press.
- Sinha, C. 2005. Blending out of the background: Play, props and staging in the material world. *Journal of Pragmatics* 37: 1537–1554.
- Sinha, C. & Rodríguez, C. 2008. Language and the signifying object. In *The Shared Mind: Perspectives on Intersubjectivity*, J. Zlatev, T. Racine, C. Sinha & E. Itkonen (eds), 357–378. Amsterdam: John Benjamins.
- Talmy, L. 1996. Fictive motion in language and “ception”. In *Language and Space*, P. Bloom, M. Peterson, L. Nadel & M. Garrett (eds), 211–276. Cambridge, MA: MIT Press.
- Tomasello, M. (ed.) 1998. *The New Psychology of Language: Cognitive and Functional Approaches to Language Structure*. Mahwah, NJ: Erlbaum.

- Tomasello, M. 1999. *The Cultural Origins of Human Cognition*. Cambridge, MA: Harvard University Press.
- Tomasello, M. 2003. *Constructing a Language: A Usage-Based Theory of Language Acquisition*. Cambridge, MA: Harvard University Press.
- Vygotsky, L. S. 1978. *Mind in Society: The Development of Higher Mental Processes*. Cambridge, MA: Harvard University Press.
- Waddington, C. H. 1975. *The Evolution of an Evolutionist*. Edinburgh: Edinburgh University Press.
- Whiten, A., Goodall, J., McGrew, W. C., Nishida, T., Reynolds, V., Sugiyama, Y., Tutin, C. E. G., Wrangham, R. W. & Boesch, C. 1999. Cultures and chimpanzees. *Nature* 399: 682–685.
- Wilson, E. O. 1998. *Consilience: The Unity of Knowledge*. New York: Knopf.
- Zlatev J. & the SEDSU Project 2006. Stages in the evolution and development of sign use (SEDSU). Paper presented to the 6th International Conference on the Evolution of Language, Rome, April.

Understanding embodiment

Psychophysiological models in traditional medical systems

Magda Altman

1. Introduction

The embodiment of cognition has become a fundamental tenet in many areas of cognitive science research. However, the nature of the body purported to be embodied is a more contentious question often influenced by specific research objectives. While the cognitive linguist is primarily concerned with the ways in which language implicitly refers to and instantiates bodily experience to communicate meaning, the expert in artificial intelligence may focus on the kind of environmental interaction a symbolic system requires in order to emulate human behaviour or learning processes. In this paper, I take the view that embodiment requires the experience of being in a body rather than just the information gathering and interactive capabilities that having a body may afford. From this perspective, language is embodied firstly, because it can provide access to body-based knowledge of perceived affordances (Gibson 1977) in the world we inhabit and secondly, because it enables communication about our inter (and even intra) subjective experiences. In short, language allows us to talk about how we experience, construe and *feel* about the world, each other and ourselves. Arguably, communicating about our subjectively felt experience is as powerful and useful an ability as being able to refer to, describe or produce propositions about 'states of affairs' in the world.

Communication about subjective realities requires a common ground of reference. We need to know (consciously or unconsciously) that when we use a term such as 'cold', the listener will infer the kind of physiological and even psychological dynamics that we identify with cold – for instance, a feeling of contraction, avoidance, perhaps even fear. It is these kinds of ontological entailments that flesh out the skeletal code of speech and make it meaningful. Below, I examine several current hypotheses concerning the role of the body in communicating representations of subjective experience. I then turn to an investigation of folk theoretical representations of the body in traditional medical psychophysiological models. My aim is to show that these models may provide important insights into conceptualisations of the body which underpin embodied cognition and language.

2. Current approaches to embodiment in the cognitive sciences

Zlatev (2007:325) proposes “to bridge (or at least minimize) the gap between language and embodiment” with ‘mimetic schemas’, which he defines as the “volitional use of the body for constructing and communicating representations” (Zlatev 2007:301). He proposes that shared ‘mimetic schemas’ allow us to ‘activate’ consciousness of the appropriate communicative intent. These mimetic schemas are derived from “imitation – which with time becomes internalized” (Zlatev 2007:323). Yet, Zlatev (2007:325) allows that each [mimetic] schema has a different emotional-proprioceptive “feel, or affective tone”, using the examples of *KISS* and *KICK*. This hypothesis presumes that because we are able to mimic (possibly virtually) the perceived motor behaviour of other humans with our own actions, we can experience the emotions or states of mind that lead to, result from or accompany these actions. Mimetic schemas, Zlatev (2007:326) surmises, are the “‘missing link’ between sensorimotor and linguistic cognition,” an early form of representation which precedes language both in evolution and ontogenesis.

However, this hypothesis still begs the question as to how, where and when embodied schematisation (as contrasted with mere observation) takes place and the possibly related issue of how an observed behaviour is mapped onto the far more complex action that we (internally) experience. Kissing and kicking involve considerably more sensory-motor as well as proprioceptive and emotional information than is presented in the visual display of a kiss or kick. Although we learn a great deal about behaviour through perception and by mimicking the actions of others, presumably we need to supply the sensory, motor and proprioceptive information which is not perceptually (for instance, visually) available. Mimesis requires a subjective internal basis – probably both a pre-perceptual understanding of action and the ability to mimic action which has been perceived.

In her research on language development, Mandler (1992:592) contests that “image schemas involve simplifying and redescribing perceptual input”. Following Lakoff (1987) and Johnson (1987), she postulates that image schemas “consist of dynamic spatial patterns that underlie the spatial relations and movements found in actual concrete images.” According to Mandler (1992:592), the ability to abstract image schemas presupposes some “innate mechanism of analysis”. This is the basis for her contention that “image schemas provide the meanings that enable infants to imitate actions.” As she concludes, this hypothesis is “more or less the opposite of that proposed by Piaget” who postulated that action is understood through a gradual process of ‘interiorization.’ Mandler does not further specify the ‘innate mechanism of analysis’ that abstracts image schemas from actions. She does suggest that “the dynamic and relational nature of image schemas provides a kind of syntax” (1992:592) which may still be pre-propositional.

A number of neuroscientists including Gallese (2005) posit that *mirror neurons* are the neural mechanism by which we understand the perceived actions of others. Gallese (2005) and Gallese and Lakoff (2005) report that observed (other) and planned (self) action share a largely common neural architecture for the simulation of action while the attribution of agency (self or other) takes place in a different neural region. Explaining cognitive processes in terms of neural substrates is a promising approach. However, critics (see Zlatev 2007) have pointed to the tendency to emphasise the brain at the expense

of the body and the complexity of phenomenological experience. Presumably, a complex interplay of biological, including but not exclusively neural, functions is involved in the subjective experience of action and all of these contribute to the nature of our *own* (imagined and actual) experiences as well as our understanding of *others'* behaviours. This kind of complexity was captured by Maturana and Varela (1980) in their notion of autopoiesis; the complexity of the living system has also been referred to as 'organismic embodiment' by Ziemke (2003: 4).

Gallagher's (2003) research also supports an internal basis for embodied schema. He argues that "If there is a form of body awareness that captures the differential spatial order of the body but does not involve an egocentric spatial framework [presumably required for the subject/other agent distinction], then that awareness would be non-perceptual" (Gallagher 2003: 57). Gallagher and Cole (1995) contrast the 'body image,' a "mental representation that one has of one's own body" (Gallagher and Cole 1995: 369) with the 'body schema' which includes 1) "input and processing of new information about posture and movement that is constantly provided by a number of sources, including proprioception," 2) "motor habits, learned movement patterns ('motor schemas' or programs)," and 3) "certain intermodal abilities that allow for communication between proprioceptive information and perceptual awareness, and an integration of sensory information and movement" (Gallagher and Cole 1995: 375–376). They conclude that "in all three of its functional aspects, the body schema system is interrelated with perceptual aspects of the body image" and thus may license, while being distinct from, conscious representation (Gallagher and Cole 1995: 376). The relationship between body schema and perceptual information as well as the manner in which the body becomes 'schematized' both require further analysis. However, Gallagher and Cole's (1995) hypothesis suggests that the body schema, as a proprioceptive and largely subconscious 'system,' could be the basis for the representation of action. If so, the subjectively experienced body becomes the initial locus of embodied cognition including the perceptual processes that would enable mimesis. As Gallagher (2003: 62) remarks, "if one accepts the premise that sense perception of the world is spatially organized by an implicit reference to our bodily position, the awareness that is the basis for that implicit reference cannot be based on perceptual awareness without the threat of infinite regress."

Gallagher's (2003) postulate that there is some system of 'communication' between proprioceptive information and sensory perception raises interesting questions in relation to research on perceptual processes. Although perception has traditionally been associated with an input-driven model and image-rich records, Barsalou (1999: 8) maintains that perceptual information is pre-processed and schematically stored: "Rather than containing an entire holistic representation of a perceptual brain state, a perceptual symbol contains only a schematic aspect." In his view, the perceptual symbol formation process works on all modalities and can provide symbols for vision, audition, haptics, olfaction, and gustation as well as for proprioception and introspection. While Barsalou (1999: 585) presumes that "each type of symbol becomes established in its respective brain area," he maintains that these symbols are not internally 'available' on a unitary basis but are related in simulations of events and event sequences. In his view, these sequences are constructed by a 'simulator' that can be used after, or in the absence of, actual events and may contain

any number of linked multi-modal perceptual symbols (e.g. haptic and visual).¹ Barsalou (2003) favours a dynamic view in which aspects of stored simulations of events are evoked, re-enacted and adapted 'online' during new encounters.

While Barsalou is concerned largely with the cognitive architecture involved in perceptual processing rather than the bodily origins of schematic structures, his work is compatible with the postulate of a more fully embodied foundation. His postulate that dynamic event schemas are an important integrating factor in 'simulations' suggests a spatial-motor framework possibly underpinned, as I suggest below, by a proprioceptive framework. In short, the body schema proposed by Gallagher and Cole (1995) may be an integral aspect of (peripersonal) spatial organization which in turn could structure Barsalou's (1999) postulated 'perceptual symbols'.

The importance of event schemas or actions in these various accounts of body-based schematics makes the force dynamics of the body an area worthy of particular attention. Although he does not specifically address a proprioceptive basis for dynamic schematics, Talmy (2000) has extensively investigated Force Dynamics (FD) in language. Talmy points out (see Manjali 1997) that the forces represented in language "continue to incorporate prescientific conceptions of interactional dynamics, which pertain more to the domain of our common sense, or to what has recently been investigated under the rubric of 'naïve physics.'" Below, I investigate the possibility that this pretheoretical understanding is represented in traditional medical models of the body. I argue that these ancient medical systems may provide one of the most insightful and complete sources of information on the subjective (and for most of us largely unconscious) experience of the body – in sum, a folk theoretical representation of the proposed body schema. Throughout the discussion, I refer to Talmy and the research discussed above, to compare and contrast modern approaches to force dynamics, image schemas and embodiment with those of traditional medical systems, in particular, Chinese Medicine.

3. Traditional psychophysiology: The case of Chinese Medicine

In my analysis of traditional models of psychophysiology, I focus on the ancient system of Chinese Medicine (CM), as the level and antiquity of Chinese textual documentation makes this tradition more accessible to research than many of the largely oral traditions. After providing some general background, I explore several key Chinese medical terms that function as probes into these 'folk' or pretheoretical conceptualizations of the body (see Figure 1).

During several historical periods in China, systematizations of CM theory were undertaken (see Unschuld 1985 and Kaptchuk in *East Asian Medical Studies Society* 1985). These systematizations were aimed at improving internal theoretical consistency, compatibility with contemporary religious, cultural and political beliefs or eradicating 'superstition' and presenting a more 'scientific' basis for the tradition. Reformulations of this

1. Support for the hypothesis that spatial/motor information is recorded separately from object features comes from work by Milner and Goodale (1995).

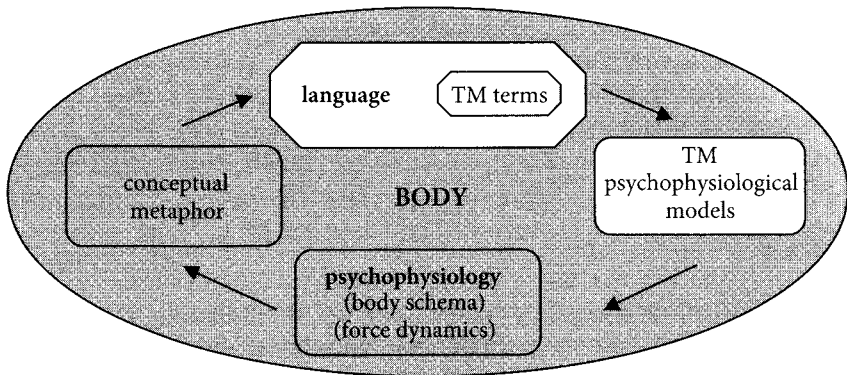


Figure 1. Traditional medical (TM) terms provide a probe into underlying conceptualizations of the body

nature often obscure underlying beliefs, and for this reason, I have tried to address only central concepts that have been consistently used since ancient times.

The Chinese models examined here are the *yin yang*, *wu xing* ('five phase') and the *jing luo* ('channel') models, which represent psychophysiological structure and function. I also consider the related concept of *qi* ('finest matter influence' (Unschuld 1985)) as *qi* is one of the most important constituents in the force dynamics of the body. These CM models are of particular interest not only for their antiquity, but also because they closely resemble those of other ancient traditional medical (TM) systems, including the Sowa Rigpa (Tibetan), Ayurveda (Indian) and Unani Tibb (Graeco–Arabic) systems. While far more research is required before any definitive conclusions can be drawn, it appears that similar models, concepts and principles are incorporated in a wide variety of TM or more broadly indigenous knowledge systems. It certainly seems possible that these not only provide culturally specific representations of the body for the purposes of medical treatment and self-development protocols but may also be based on a more universal body schema.

It is perhaps not surprising to discover a level of universality in basic Chinese medical thinking given the fact that most of these medical principles were presumably established through intensive study of how the body *feels*. TM concepts and models rely heavily on the subjective experience of the patient and the subjective evaluation of the physician who – with only minimal technological assistance – is required to assess complex normal and pathological psychophysiological states and their likely progressions using his/her own perceptual and intuitive capacities. Most TM systems also include a variety of practices for adepts, aimed at improving the function of the body–mind. These practices entail becoming progressively more aware of (and ultimately learning to enhance or transform) otherwise unconscious psychophysiological forces and structures and can be considered further sources of introspective knowledge about bodily processes (see Jou 1980).

The oldest available written records on CM include the inscribed tortoise shells used for divination, some of which date back to as early as the 16th century BCE.² The first known theoretical texts are the *Ma wang dui* silk manuscripts from 168 BCE (Bai 2001: 10) and the somewhat later, more comprehensive *Huang di nei jing* (Unschuld 2003: 76).³ The concepts discussed below were already present in one or both of these early texts.

3.1 *Yin, yang and qi*

In his work on Daoism, Schipper (1993) explains that the primordial chaos which precedes creation is called *hun tun*, an onomatopoeic expression for the chaotic “sphere or matrix that holds within itself the whole universe, but in a diffuse, undifferentiated and potential state” (Schipper 1993: 34). The *hun tun* contains the original *qi* (‘breaths’) before the moment of creation. At the moment of creation these *qi* escape: the “light, transparent *ch’i* [*qi*] rise and form heaven; the heavy, opaque ones sink, forming Earth. Having established the polarity of Heaven and Earth, the *c’hi* [*qi*] join and unite in the Center, which constitutes a third fundamental modality [the human]” (Schipper 1993: 34). This original primordial and undifferentiated state is also called *wu ji*. The Oxford Concise Chinese–English Dictionary (Manser 1999) translates *wu* as ‘nothing, nil’ and *ji* as ‘the utmost point, extreme, pole’. *Ji* can also be translated as ‘polarity’ making *wu ji* a state of non-polarity. The *tai ji* (‘great polarity’, ‘supreme, ultimate’) symbol thus represents creation as the emergence of the *yin* and *yang* polarities, their interdependence, interactions and transformations (Jou 1980) (see Figure 2).



Figure 2. The *tai ji* depicts the dynamic interaction of *yin* (black areas) and *yang* (white areas)

2. Tortoise shells were subjected to heat and fine cracks developed on the shell; the pattern of cracks was interpreted for the prediction and prognosis of matters relating to human health, the determination of necessary sacrificial rites and other important sociocultural activities. This emphasis on the significance of patterns is an essential feature of traditional CM (Kaptchuk 1983). Information on the questions addressed in the divination, the prognosis and actual outcomes was inscribed on the shell, providing a kind of clinical record.

3. The *Huang di nei jing* was compiled from the 1st and 2nd centuries BCE to the 2nd century CE (Unschuld 2003: 76) but amended, annotated and revised in later centuries. It is widely considered the most important text of CM and includes most of the basic concepts still presented in modern theoretical and practical texts.

The Chinese concept of *yin yang* dates back several millennia and provides an interesting insight into the 'naive' understanding of physics that Talmy posits must underpin the force dynamics of language (Manjali, 1997). Talmy (2000: 6) proposes that "underlying all more complex force-dynamic patterns is the steady-state opposition of two forces". In this folk view, he notes, each entity will have an *intrinsic* tendency towards action or rest. This passage from the ancient *Huang di nei jing tai su* describes exactly this kind of 'steady-state opposition' in which *yin* and *yang* are the two forces exhibiting intrinsic tendencies towards rest or action:

Yin and *yang* are the underlying principle of heaven and earth; they are the web that holds all ten thousand things secure; they are father and mother to all transformations and alterations [...]. Heaven arose out of the accumulation of *yang*; the earth arose out of the accumulation of *yin*. *Yin* is tranquillity, *yang* is agitation; *yang* creates, *yin* stimulates development; *yang* kills, *yin* stores. *Yang* transforms influences, *yin* completes form. When cold [*yin*] reaches its zenith, it creates heat [*yang*]; if heat reaches a zenith, it creates cold. (Unschuld 1985: 283–284)

In sum, *yin yang zhi yi* ('*yin* and *yang* counterbalance each other') (Wiseman and Ye 2000: 708). Yet, the *yin yang* model differs from Talmy's (2000) description of force dynamics in predicting a continual dynamic cycle of interactions from the quiescent, potential state of *yin* to the active, kinetic phase of *yang* and back. Moreover, in the *tai ji*, the next phase is always incipient in the current state as indicated by the contrasting dot of *yin* within the *yang* area and vice versa. More research would be required to apply the alternating dynamic proposed by the *tai ji* model to the language data Talmy (2000) has addressed, but the reiterative alternation between state, action and new state appears to be a fairly basic aspect of many grammatical and discourse constructions.

In Talmy's (2000) view, the FD of language are qualitative: forces have manner and are conceptualised relationally rather than quantitatively. Either the 'agonist' (the subject) or the 'antagonist' (the force (re)acting on/to the subject) is presented as the stronger or weaker in any interaction and may tend either towards rest or towards action. In the CM *tai ji* model, any interaction can be understood as a dynamic in which either *yin* (restive) or *yang* (active) dominates. The profile of any entity or agent in terms of its *yin yang* characteristics is determined by its qualities. Slow, restive, potential, dark, hidden, dense, contractive, condensing and structured are *yin* qualities while quick, active, actualized, bright, apparent, subtle, expansive, dispersive and transformative qualities are *yang*. Because the *tai ji* model provides for a qualitative categorization of *yin* and *yang*, it allows for a predictive model in which agents and entities interact and event dynamics unfold in accordance with their relative *yin yang* profile. Applied to a prototypical sentence in which an agent (*yang*) acts (*yang*) to produce a change (*yang*) of state (*yin*) in the theme, this provides a similar analysis to the trajectory-landmark schematic postulated by Langacker (2000). Langacker posits that an agent and theme are distinguished in terms of energy transfer: the agent is the source of energy and the theme the terminus of the action chain, the ultimate "energy sink" (Langacker 2000: 30).

In Talmy's view, entities/agents not only exert forces but can in other ways influence and alter each others' tendencies through, for example, resistance to force, overcoming of

resistance, blockage to force exertion, and removal of blockage (Manjali 1997). This brings us to a more detailed discussion of the force dynamics of the body in CM, more specifically to a consideration of the flow of *qi* through the *jing luo* ('channels') of the body.

3.2 Dynamics of *qi* and the *jing luo*

The notion of *qi* has its roots in a divinatory view of the universe in which forces are enspirited and ubiquitous but not specifically human or even biological in nature. In the Chinese, as in most ancient TM systems, the practice of 'medicine' originally had a far broader scope, encompassing not only the health of the body and mind but also the well-being of the natural, social and even cosmological whole (Unschuld 1985). Natural forces, particularly *feng* ('winds,' 'climates') were closely aligned, perhaps even identified, with the intentions of ancestors who played an important role as deceased but actively concerned (and sometimes offended) family members, as intermediaries between human life and ecocosmic forces and in some cases (particularly in the later violent 'Warring states' period before the unification of China) as troubled and troublesome *gui* ('ghosts') stuck in the earthly realm. While divinatory practices evolved and to some extent were replaced by systematic doctrines, elements of these beliefs were preserved or adapted in later knowledge systems. Unschuld (1985) has suggested that *feng* is conceptually related to subsequent notions that played an important role in theories of health and disease such as *gui* ('ghost,' 'spirit') and, most importantly for this discussion, *qi*.⁴

Unschuld's (1985) translation of *qi* as the 'finest matter influence' expresses the dual material and functional, almost animate, aspects of *qi* quite precisely. Noted CM historian Sivin (1987:47) quips: *qi* is "stuff that makes things happen", "what makes things happen in stuff", and "stuff in which things happen". *Qi* is everywhere in CM ecocosmology and psychophysiology: as the original 'breaths' it inspires the creation of the universe; in the environment, it manifests as climates and winds. Physiologically, *qi* is generated as the 'kidney *yang*' (fire) heats the 'kidney *yin*' (water) causing *qi* ('steam,' 'gas') to rise and circulate through the body (see Maciocia 1989).⁵ The physiological cycle of *qi* can be likened to a water cycle (Altman 2004) in which heat (*yang*) causes a fine mist to rise and circulate while cooling (*yin*) causes condensation (see Figure 3).⁶

4. The extraordinary polysemy and cultural specificity of the term *qi* has led many modern translators of CM texts to borrow this term from the Chinese rather than attempt a translation.

5. This original *qi* is supplemented by other types of *qi* which are received and transformed by the body such as the *gu qi* (nourishment from 'grains' or foodstuffs) and the *kong qi* (from air).

6. In CM the 'kidneys' play an important role in the physiology of the body including 'governing water,' storing the inherited 'essence' (genetic function), regulating sexual vitality and reproduction and more generally constituting the 'root,' or basic vitality of the body (see Wiseman and Ye 2000:324).




Figure 3. The physiological and environmental circulation of *qi* as a ‘water cycle’

3.3 The *jing luo*

A key factor in describing (patho)physiological patterns in CM, is determining the proper circulation and flow of *qi* (and its *yin* counterpart, *xue* (‘blood’)) through the numerous ‘channels’ in the body. The theory of the *mai* (‘conduits’, ‘vessels’, ‘pathways’, ‘river branches’) (see Wiseman and Ye 2000: 654) is central to CM physiology. Many scholars of CM believed that the *acumoxa xue* (literally ‘holes’ or ‘caves’ but generally termed ‘acupuncture points’ in the English literature) predated the theory of the *mai*, which were often described as purely theoretical constructs devised to connect the *xue*.⁷ However, the *Ma wang dui* texts discovered in 1973 mention only the *mai* suggesting that the conceptualization of these channels is very ancient and may pre-date that of the *xue* (Bai 2001: 9–10). In later texts, the *mai* are also referred to as the *jing mai*, *jing*, *jing luo* or *jing mai* where *jing* denotes a ‘conduit’, ‘channel’ or ‘warp’ (in a woven tapestry) and *luo* means ‘to tie up’, ‘to fasten’, ‘to catch in a net’ (Porkert and Ulman 1982: 129) or ‘to twine, to connect’ (Bai 2001: 56). The *mai* or *jing* are largely vertical channels intersected by the *luo* to create a network of *qi* moving in a structured fashion throughout the body. A similar theory in Ayurvedic and Tantra Yoga physiology refers to the *ojas* (‘finest essence’, ‘vital energy’), *nadi* (‘conduits’) constituting the *tantra* (‘web’) of the *sukshma-sharira* (‘subtle body’) (Wujastyk 2001; Hartzell 1997).

While *qi* flowing through the network of *jing luo* describes the space of the body, the transmission of *qi* also plays a role in perpetuating life across time. The *yuan qi* (‘original *qi*’) is received from the parents at birth and maintains the body through life, being slowly depleted by age and excess or inappropriate activity. The polysemy of *jing* (a polysemy shared with the Sanskrit term *tantra*) suggests that conceptually the notion of a ‘warp’ or ‘conduit’ also has a temporal sense. The terms *jing* and *tantra* denote a literary classic, a book “recording the Way” (Zhang shi lei jing in Unschuld 2003: 15) which underscores the importance of the transmission of knowledge across time. Unschuld (2003: 16) writes: “If

7. The term ‘acumoxa’ is preferred over the more common translation ‘acupuncture’ as the Chinese term *zhen jiu* means ‘needle moxa’ (Hsu 1999). Needling refers to acupuncture, while moxa refers to burning herbs (usually *artemesia annua*), on the points.

society is comparable to a fabric combining many threads of ideas and levels of hierarchy, a certain wisdom may be considered the warp holding it all together and ensuring its everlasting functioning.” The role of the shaman in early divinatory practices was precisely to preserve the continuity of this fabric through the dimensions of time and space through contact with the ancestors.⁸ The modern CM physician is more concerned with maintaining the integrity of the pattern in the body and its harmony with natural patterns of *qi* in the environment. There is a homology between the structure of the natural and social worlds and that of the body.

In accordance with *yin yang* theory, the *jing* (‘conduits’) are identified as *yin* when the direction of flow is upward from earth to heaven or as *yang* where the flow goes from heaven to earth. The body itself is conceptualized as *qiu* (‘ball’, ‘sphere’). The periphery is defended by the *wei* (‘outer’, ‘defensive’) *qi* and the inner body nourished by the *ying* (‘camp’, ‘army’) *qi*.⁹ The *qi* flows through all the 12 main channels in an iterative, cyclical pattern, also reinforced by daily and annual rhythms in which certain channels experience increased flow. Within the network of the *jing luo*, *qi* pools and collects at the *xue* (caves, holes, acupoints). Porkert and Ulman (1982) have proposed that *yin* can be described as ‘structive’ and *yang* as ‘active’ and that any system can be viewed as a dynamic having more structive and active parts. For instance, in CM physiology, the channels are more active (and more superficial) than the organs or viscera. At the level of the organs, the *fu* (*yang* viscera) play a more active (transforming) role, while the *zang* (*yin* viscera) play a more structive (refining and storing) role. At the level of the channels, the *xue* where the *qi* collects are the more structive aspect of the flowing *jing* (‘conduits’, ‘vessels’). (From a scientific viewpoint, the *xue* are points of reduced electrical activity.) In this view, the body is a dynamic system, and the system as a whole as well as all of its parts are characterized by an interaction between structive (or stative) and transformative (or active) aspects. The smooth flowing of *qi* in the *jing luo* results in a state of health while blockage, incorrect directional flow, lack or excess of quantity or flow speed in the *jing* results in disease. While there are many other important aspects to CM psychophysiology, including the skin, muscles, bones, viscera and a variety of other structures and substances, the psychophysiology of the body can be schematically described in the circulation of *qi* through the viscera and the larger network of *jing luo* constituting the body.

3.4 CM body schema(s)

The complex circulation of *qi* in the body according to the principles of *yin yang* appears central to the representation of the body in CM. The dynamic pattern of force interactions

8. This role is still played by many shamans or traditional healers worldwide. In South Africa, one of the important roles of the *izangoma* (‘diviners’) is to help maintain respectful and beneficial relations between the living and their *amadlozi* (‘ancestors’).

9. The militaristic metaphors underpinning a number of CM terms including *wei* and *ying qi* probably date back to the violent and divisive ‘Warring states’ period before China became a single unified nation (Unschuld 1985).

involved is also essential to the CM body schema. All of the aspects of force mentioned by Talmy in his discussion of the FD of language also pertain to the flow of *qi* in the *jing*. *Qi* normally flows in a structured fashion in the *jing* and between the viscera but the normal manner and direction of flow may be disturbed in various ways, for instance, blocked, dispersed, reversed, depressed, invaded or subdued. This is reflected in diagnoses such as:

- a. *gan qi fan pi* ('liver *qi* invades spleen') (Wiseman and Ye 2000: 361);
- b. *wei qi ni* ('stomach *qi* counterflow') (Wiseman and Ye 2000: 102);
- c. *shen bu na qi* ('kidneys fail to grasp *qi*') (Wiseman and Ye 2000: 326);
- d. *fei qi bu xuan* ('lung *qi* doesn't diffuse') (Wiseman and Ye 2000: 413); or
- e. *zao xie fan fei* ('dry evil invades the lungs') (Wiseman and Ye 2000: 372).

Similar aspects of force dynamics were highlighted by Johnson (1987) who postulated that FORCE, BLOCKAGE, REMOVAL OF RESTRAINT, ENABLEMENT, DIVERSION and COUNTER FORCE were among the most important image schemas based on the body.

One of the most widely discussed schemas in the cognitive linguistic literature is the PATH schema. In CM, the *jing* are clearly important as the pathways of the living body and are involved in all psychophysiological processes. The body dynamics described by the flow of *qi* through these pathways embody many of the other schemas described by Johnson (1987) and other researchers. UPWARD and OUTWARD schemas are instantiated as the *qi* rises upward and disperses outward; DOWNWARD and INWARD movement occurs as the *qi* descends to the inner and lower *yin* areas in the descending aspect of the cycle of *qi*. A difference I have underscored by adding the 'ward' suffix to UP, DOWN, IN and OUT is that all schemas in the CM system are intrinsically dynamic. Their basis in the circulation of *qi* means they are also CYCLICAL and ITERATIVE – two other important image schemas noted by Johnson (1987). The fundamental FORCE and COUNTERFORCE schemas are embodied in the flow of *qi* through the *jing*. The CONTAINER and PERIPHERY schemas are represented by the 'ball' of *qi* comprising the *nei* (INSIDE) *qi* of the inner body and the *wei* (OUTSIDE) *qi* that guards the periphery of the body. In short, the dynamics of *qi* flowing through the *jing luo* could constitute a proprioceptive basis for a body schema which includes the various images schemas noted by Johnson (1987) and others.

While a bodily basis for image schemas is not in itself a new argument (see early work by Johnson 1987, and Lakoff 1987), the proposal here differs in suggesting that image schemas such as IN, OUT, UP and DOWN may initially be based on proprioception. This means that they are subjectively experienced in an integrated fashion in the force dynamics of the body rather than objectively observed on a one-by-one basis in the world and later combined to produce more complex schematics. A typical CM diagnosis describes a pathology in terms of a patterns such as those listed in examples a–e above, which only make sense in terms of the larger model of CM body dynamics or 'energetics'. The integrated nature of CM schematics is also apparent with respect to the many polar schemas investigated by Johnson (1987) and others. Although Johnson (1987) noted that the CYCLE is an important schema, he did not relate this to polar schemas. In his view and in much of the cognitive linguistic literature, polar schemas are statically represented on a linear continuum. However, in the CM models, up/down, in/out – as well as other polarities less frequently mentioned in the cognitive linguistic literature such as light/dark and hot/

cold – are identified with the *yin* (downward, inward) and *yang* (upward, outward) aspects of the dynamic *tai ji* cycle.

The cyclical CM model also helps make sense of the kinds of complex relationships observed in the semantics of polarities or axiology. Krzeszowski (1993) has proposed that there is an axiological parameter implicit in polar schemas; for instance, UP is good and DOWN is bad. A difficulty with this proposal is that UP sometimes has a negative and DOWN a positive connotation. This is seen in the negative implications of many English colloquial expressions using UP such as ‘getting uppity’, ‘having one’s head in the clouds’ and being ‘over one’s head’. The positive aspect of DOWN is evidenced in English expressions such as ‘well-grounded’, ‘down to earth’, ‘getting down to basics’, ‘well-founded’ and the more slang phrasal verb idioms ‘to get it down’ (to fully understand), ‘to get down’ (to have a good time) and ‘to be down with’ (to agree). This variation in axiological semantics makes sense in terms of the cyclical schematic represented by the *tai ji*: the value attributed to any directional movement is determined with reference to the need for an overall balance of *yin* and *yang*. In other words, the steady state or homeostatic imperatives constrain the dynamic transformations and relative strengths of *yin* and *yang* in any process. While UP is generally associated with the positive *yang* attributes of warmth, joy and consciousness, excessive upward movement, (*yang* rising without grounding in *yin*) is undesirable; the grounding foundational and ‘structutive’ aspect of *yin* is essential in body dynamics. The cyclical and recursive dynamic of the *tai ji* model may have some advantages over the linear continuum model in allowing for the otherwise difficult to explain variation in the connotations of polar or axiological terms.

Talmy (2000:422) writes that: “All of the interrelated factors in any force-dynamic pattern are necessarily co-present wherever that pattern is involved. But a sentence expressing that pattern can pick out different subsets of the factors for explicit reference – leaving the remainder unmentioned – and to these factors it can assign different syntactic roles with alternative constructions.” In CM, the term *pattern* is widely used in assessing body dynamics and foregrounds those aspects of the circulation of *qi* and *yin yang* parameters (for instance, the levels of cold/heat, the inner/superficial location of disease and the overall ‘vacuity’/‘repletion’) that are most significant in identifying a condition (see Wiseman and Ye 2000:169). However, while a particular pattern such as ‘deficient *yin*’ or ‘excess heat’ may be highlighted, all of the interrelated factors in the overall pattern of body dynamics are necessarily involved. Talmy notes that the FD of language allow this type of selective highlighting so that any chain of interaction is subject to both paradigmatic and syntagmatic reduction (Manjali 1997). Paradigmatic reduction allows for properties and conditions rather than agents or entities to have effects and is evidenced in CM diagnoses such as ‘heat damages heart *yin*’ or ‘*yin* deficiency causes empty heat to rise’. CM descriptions of the (patho)physiological states of *qi* in the body also make use of syntagmatic reduction from a larger causal continuum. Maxims such as ‘the liver *qi* invades the spleen’ or ‘lungs fail to descend and kidneys fail to grasp *qi*’ are relationships picked out from the larger pattern of interactions described by the *wu xing* (‘five phase’) model (see discussion below).

Mandler’s (1992:592) notion that image schemas “consist of dynamic spatial patterns that underlie the spatial relations and movements found in actual concrete images” has been noted above. She postulates that “image schemas provide the meanings that enable

infants to imitate actions” and that “the dynamic and relational nature of image schemas provides a kind of syntax” (Mandler 1992: 592). It is certainly interesting that the CM psychophysiological model provides an integrated presentation of these same dynamic spatial patterns. Notably, however, the CM image schemas are initially derived from proprioception, the subjective experience of the body, rather than from sensory-motor perceptual experience in and of the world. Correlating ancient representations of the subjectively experienced body with accounts of Force Dynamics in language (Talmy 2000; nd), research on the image schematic bases of language (Johnson 1987) and scientific information on neural structures representing action in the brain (see Gallese and Lakoff 2005) is clearly a tall order. Nevertheless, a proprioceptive basis for cognitive processes is precisely what an embodied view of language points to, making further research worthwhile.

3.5 The *wu xing*

Hsü Ta-ch'un (in Unschuld 1998: 9) wrote: “Each of the five viscera has its genuine essence; these [essences] are separate materializations of the original influence [inherited at birth]”. In the cycle of *qi*, the activities of each viscus are associated with different qualities such as heat, cold, damp, dryness and wind. Each has a facilitating or restraining effect on the other viscera. Together, the viscera support the dynamic transformations of the body and keep it tuned to the changing environmental conditions and seasonal cycles with which they correspond. In this way, human physiology reflects the balance of forces found in nature. As written in the *Huang di nei jing* (anon. translation in Hsu 1999: 108):

Heaven has Five Phases, and they couple with the Five Directions, thus giving birth to the Cold, the Hot, the Dry, the Damp, and the Windy. Man has Five Organs, and they change the Five *qi*, thus giving birth to Joy, Anger, Worry, Sorrow and Fear. The ‘Discussion’ says that the Five Cycles [of the Five Phases] mutually ride each other and that they all govern each other and that on the day when one arrives at the end of the year the circle is complete and starts all over again...

The *wu xing* or ‘five phases’ referred to above are generally identified as fire, water, wood (tree), soil and metal.

Xing is a term which is variously translated as ‘step’, ‘element’, ‘(evolutive) phase’ (Porkert and Ulman 1982) and ‘agent’ (Unschuld 2003).¹⁰ As with *yin*, *yang* and *qi*, the nature of the *xing* is partly substance, partly quality, partly activity. The *xing* have to interact properly, both supporting and restraining each other. Consequently, numerous supportive and destructive ‘cycles’ are described in the CM literature.

Interestingly, many other systems of TM have their own versions of these *xing*. Ayurveda and Sowa Rigpa (Tibetan medicine) respectively have *prthivii* or *sa* (‘earth’), *aap* or *chu* (‘water’), *tejas* or *me* (‘fire’), *vaayu* or *rLung* (‘wind’) and *aakaasha* or *nam-mkha*

10. Unschuld (2003: 84) also notes that *xing* is closely related to the ancient concept of *de* (‘virtue’, ‘power’) a concept which is central to the ancient philosophical text attributed to the sage Lao Zi, the *Dao de jing* (‘Way [and its] power classic’). “Both *de* and *xing* appear to have conveyed an identical meaning of ‘virtues’ or ‘agents’ in the sense of forces that make certain things or processes happen”.

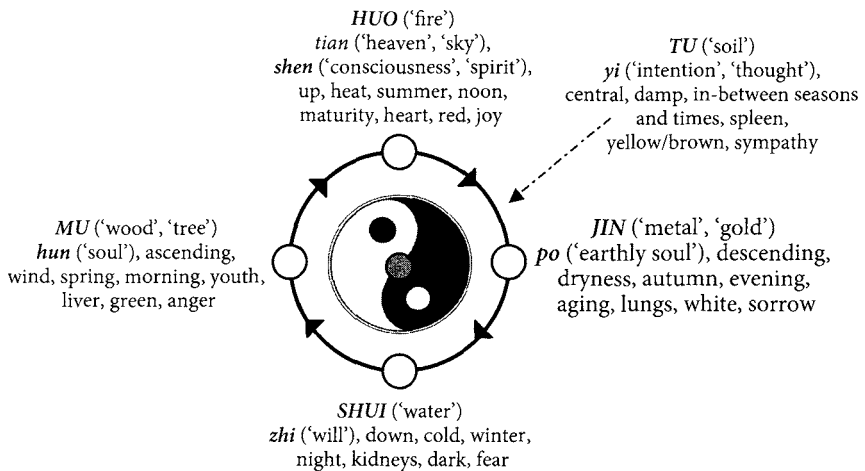


Figure 4. The *wu xing*: the *sheng* ('engendering') cycle in which water nourishes wood, wood feeds fire, etc.

('what receives the light' or 'space') (Men-Tsee Khang 2001). Unani-Tibb (Graeco-Arabic medicine) has only four 'elements,' earth, water, fire and air, but in both the Ayurvedic (Wujastyk 2001) and Unani-Tibb (Bikha and Haq 2001) systems, the various 'elements' are considered to combine to constitute 'humours,' 'qualitative tendencies' or 'influences.' In all these systems, it is ultimately the appropriate location, equilibrium and correct transformations of the various influences as they move through the viscera and conduits that guarantee or constitute health.

Before considering the qualitative aspect of the *xing*, it is worthwhile looking at the *xing* in terms of the force dynamics and neural architecture discussed above. Porkert and Ulman (1982: 74–75) use the metaphor of the felling of a stag in a hunt to describe the "shifts of polarity" in the cycle of "potentiality and actuality" represented by the *xing*: the wood *xing* (in their terms "transformative phase") of potential activity corresponds to the drawing of the crossbow; the fire phase to the actual activity as the bolt is in flight; the "potentially constructive" phase of metal to the moment when the bolt has struck its quarry but the stag is still in flight; and the final "actually constructive" water phase to the endpoint when the stag is dead. In Porkert and Ulman's view (1982: 75), the earth phase is the period in which the "phase shifts" take place, the decisive in-between moments in which it is "not precisely apparent what the next phase is going to be."¹¹

The phase structure attributed to the *wu xing* cycle provides an interesting parallel with the analysis of action schemas in the cognitive linguistic and neuroscience literature. Gallese and Lakoff (2005: 467) state that: "Schemas are interactional, arising from (1) the nature of our bodies, (2) the nature of our brains, and (3) the nature of our social and physical interactions in the world. Schemas are therefore not purely internal, nor are

11. Porkert and Ulman's (1982) theory of the earth phase does accord with many versions of CM *xing* theory in which the earth *xing* also represents a short in-between 'season' of consolidation between the other four seasons in the yearly seasonal cycle. Also see Maciocia (1989).

they purely representations of external reality.” According to Gallese and Lakoff (2005), a schema is a dynamic pattern of action which underpins what is called a ‘concept.’ Using ‘grasp’ as an example, they specify the various parameters which neurologically constitute ‘grasp’, for instance roles (agent, object, object location and action), phases (initial condition, starting phase, central phase, purpose condition, ending phase and final state) as well as manner and parameter values. While phase dynamics can be attributed to any system, this parallel is noted because the first, the CM *wu xing* model, pertains to subjectively experienced or proprioceptive body dynamics while the second is purported to account for neural structures (functional clusters of neurons) that structure both motor and perceived motor events. In other words, the CM model not only incorporates image schematics but also the phase structure of event dynamics. The pause between phases that Porkert attributes to the earth *xing* has its counterpart in the middle of the seven-phase ‘pre-motor’ neural model developed by Narayanan (in Gallese and Lakoff 2005). This model includes the computational equivalent of the ‘earth’ reflective phase in which an action is checked, and possibly discontinued or reiterated.

3.6 Multimodality

Another noteworthy aspect of the *wu xing* model is that the force dynamics of the body are associated with sensorial, emotive, mental and other qualities and therefore provide a basis for (cross) categorisation. This invites further investigation particularly in light of the interest in multimodal cognitive structures in recent cognitive neuroscience research. Above, we noted Gallagher and Cole’s (1995) description of a body schema as involving “a system of motor capacities, abilities, and habits that enable movement and the maintenance of posture” (Gallagher and Cole 1995: 371) and “certain intermodal abilities that allow for communication between proprioceptive information and perceptual awareness, and an integration of sensory information and movement” (Gallagher and Cole 1995: 376). I have suggested that CM models may be based on this type of proprioceptive information about the body – although the dynamic structures they refer to are not only motor but also include what is sometimes termed in the CM literature ‘energetic’ (i.e. force dynamical) patterns and schemas. The *wu xing* model additionally provides a unified theory in which haptic and other sensorial experience is correlated to force dynamics. These correlations go further in including the psychological aspects of subjective experience.

Barsalou (1999) maintains that perceptual symbols from a range of sensory and even proprioceptive modalities are combined in *simulations*. As he writes, “During the recollection of a perceived object [...] conjunctive neurons re-enact the sensory-motor and introspective states that were active while processing it originally” (Barsalou 2003: 1180). The *wu xing* model describes an embodied basis for the type of ‘conjunction’ that Barsalou identifies at the neural level. However, the *wu xing* model goes further in providing a schematic which links certain categories of experience across domains. For instance, fear, cold, dark and winter are correlated with the contractive, descending *yin* aspect of the *tai ji* cycle and the water *xing*. In other words, these correlations are psychophysiological rather than just associated on the basis of particular experiences. At

the same time, since the body reflects or, more precisely, is homologous with the natural world, the model predicts that experiences in the environment are likely to exhibit the same kind of correspondences. The bodily basis (as contrasted with a purely experiential or environmental basis) for many of the associations formalized in the *wu xing* categories deserves more extensive and systematic investigation. However, a cursory glance at common idioms in English such as 'frozen with fear', 'enlightened', 'warm feelings', 'dark mood' appears to confirm the psychophysiological basis of many of these cross-modal and cross-domain associations.

At the beginning of this paper, I noted that to use language effectively we require some way of knowing that the term *cold* evokes not only knowledge of a temperature range, but also familiarity with the diversity of psychophysiological experiences cold entails. In the *wu xing* model, the term cold has exactly these ontological entailments.

3.7 Conceptual metaphors

Numerous cognitive linguists (Lakoff 1987; Johnson 1987) have conducted research on the metaphorical extensions made possible by underlying schemas. The schema UP, for instance, is considered to underpin metaphorical representations of happy emotions, good social standing, and, more generally, increase and improvement in a wide number of domains. We feel 'up', wake 'up', move 'up' in our careers, heat 'up', turn 'up' the lights and so on. Notably, many of these uses of 'up' are associated with an upward directional dynamic PATH rather than just a relative static position so that the term UPWARD might be more appropriate for this schema. These multimodal and multi-domain associations with UP(WARD) make sense in terms of the *wu xing* and underpinning *yin yang* models. The movement of *qi* as it passes through the heart *xing* in the zenith of the cycle is experienced as joy, increased alertness and a light spirit. (In CM as in many other TM systems, the heart is considered a seat of consciousness; it is said to house the *shen* ('spirit'); it is also considered to occupy the highest position in and to 'command' the body). The CM body schema appears to offer considerable insight into common idioms and metaphors purported to be based on 'folk' or 'pre-theoretical' models of the body (see Altman 2004 for further examples).

The complex schematic dynamics and multimodal integration described by the various CM models also provide insight into numerous conceptual metaphors discussed in the literature. An interesting example is the conceptual metaphor ANGER IS THE HEAT OF A FLUID IN A CONTAINER (Lakoff 1987; Kövecses 1986, 2000). This metaphor has been analysed in terms of underlying image schemas: the body is conceptualised as a CONTAINER unable or barely able to contain the increased and rising FORCE of the heated (angry) fluid (emotion). The CM model elegantly explains the blending of schemas in this conceptual metaphor. The excessive surge of *qi* in the upward *yang* (wood) aspect of the body cycle is experienced emotionally as anger. The *qi* moves upward from the wood *xing* (correlated with both wind and anger) into the fire *xing* where it manifests as heat, a red face, excitability and agitation. The term *sheng qi* ('rising' *qi*) means 'anger' in Mandarin, further confirming the psychophysiological equivalence of heated fluids rising in the body (or,

in the CM model, excessive rising *qi* ('steam')) with anger. It certainly seems possible that this metaphor is based on proprioceptive experience rather than the mapping of emotion onto abstract image schemas derived from the observation of physical objects and forces in the world. In other words, we do not need to discover that fire heating water in a closed container produces steam and potentially makes that container explode. We experience this process in the body container as 'fire' (*yang*) heats the body 'water' (*yin*) to produce *qi* ('steam'). If this *qi* is not cooled and condensed by the appropriate downward flow in the cycle (as it moves into the cooling and condensing metal *xing*), it will cause the extreme mental and physical agitation we call anger.

Traditional medical systems such as CM have the advantage of being holistic, in the sense that they do not merely offer piecemeal insights into isolated conceptualizations of the body such as the BODY IS A CONTAINER or EMOTIONS ARE FORCES (Lakoff 1987) but comprehensive models in which the experience of our anatomy (often rather mechanistically referred to as 'body parts' in the cognitive semantic literature), physiology and even psychology is integrated in a force dynamical schematic of the body. The example of ANGER IS THE HEAT OF A FLUID IN A CONTAINER shows the value of examining metaphors based on a 'pretheoretical' or 'naïve' understanding of force dynamics in terms of traditional medical models.

3.8 The body and the world: CM homologies

A major challenge in cognitive linguistic theory is that the 'body' postulated to be the basis for embodiment is largely implicit and often pre- or subconscious (Lakoff and Johnson 1999) making it largely invisible to language users. The CM models briefly examined in this paper offer explicit representations of the subjectively experienced body making it more accessible to analysis. Another difficult issue in linguistics is the 'symbol grounding problem' (Harnad 1990): the incompatibility of an arbitrary, abstract symbolic system with embodied meaning. In the CM model, the body is the locus of our physiological, psychological and multimodal experience and thus provides an ontological and epistemological basis for meaningful experience in the world. CM, like other traditional medical systems, presumes that the body (psychophysiology) and the natural world (ecocosmology) are homologous. This homology provides the basis for the correlations between the structure and force dynamics of the body and those of nature providing a folk theoretical explanation for the ways in which reference to the 'world' is enabled by reference to the body and vice versa (see Kaptchuk 1983 on the CM interpretation of the body as a landscape). Using current terminology, I am suggesting that the proprioceptive experience of the force dynamics of the body as represented in CM models may constitute the body schema which is the basis for the action and event schemas that structure our experience in the world and their multimodal correlations.

4. Conclusion

In this paper, I have argued that the psychophysiology of traditional medical models may offer an excellent source for insight into 'naïve', 'folk' or 'pre-theoretical' views of the body. TM models remain a largely neglected source of information on how our ancestors (and large numbers of our contemporaries) conceptualise the body and its relation to the natural and social environments. My aim has not been to scientifically validate or invalidate traditional medical perspectives on the body but rather to explore them as cohesive and well-elaborated representations of the (predominantly subjective) experience of being in a body. To this end, I have very briefly analysed a number of important models in Chinese medicine, an admittedly cursory examination of a vast and complex tradition. This preliminary investigation aims merely to demonstrate that this and the many other TM psychophysiological models merit further research and to suggest that these ancient representations of the body may complement scientific investigation in helping us understand the nature of embodiment and the kind of 'body' language instantiates.

References

- Altman, M. 2004. *Making Sense of Traditional Chinese Medicine: A Cognitive Semantic Approach*. MA thesis, University of South Africa.
- Bai, X. 2001. *Acupuncture*. Oxford: Butterworth.
- Barsalou, L. W. 1999. Perceptual symbol systems. *Behavioral and Brain Sciences* 22 (4): 577–660.
- Barsalou, L. W. 2003. Abstraction in perceptual symbol systems. *Philosophical Transactions of The Royal Society of London: Biological Sciences* 358: 1177–1187.
- Bikha, R. & Haq, M. A. 2001. *Tibb*. Gauteng: Mountain of Light.
- East Asian Medical Studies Society. 1985. *Fundamentals of Chinese Medicine*, Zmiewski, P. (ed.). Brookline: Paradigm Publications.
- Gallagher, S. 2003. Bodily self-awareness and object perception. *Theoria et Historia Scientiarum: International Journal for Interdisciplinary Studies* 7 (1): 53–68.
- Gallagher, S. & Cole, J. 1995. Body schema and body image in a deafferented subject. *Journal of Mind and Behaviour* 16: 369–390.
- Gallese, V. 2005. Intentional attunement. The mirror neuron system and its role in interpersonal relations. <http://www.interdisciplines.org/mirror/papers/1>.
- Gallese, V. & Lakoff, G. 2005. The brain's concepts: The role of the sensory-motor system in conceptual knowledge. *Cognitive Neuropsychology* 21 (0): 455–479.
- Gibson, J. J. 1977. The theory of affordances. In *Perceiving, Acting and Knowing*, R. Shaw & J. Bransford (eds). Hillsdale, NJ: Erlbaum.
- Harnad, S. 1990. The symbol grounding problem. *Physica* 42: 335–346.
- Hartzell, J. F. 1997. *Tantric Yoga*. PhD Dissertation, Columbia University.
- Hsu, E. 1999. *The Transmission of Chinese Medicine*. Cambridge: Cambridge University Press.
- Kaptchuk, T. J. 1985. Introduction. In *Fundamentals of Chinese Medicine*, East Asian Medical Studies Society (ed.), xvii–xxxvi. Brookline: Paradigm Publications.
- Johnson, M. 1987. *The Body in the Mind: The Bodily Basis of Meaning, Imagination, and Reason*. Chicago: University of Chicago Press.
- Jou, T. H. 1980. *The Tao of Tai-Chi Chuan*. Taiwan.

- Kaptchuk, T. J. 1983. *The Web That Has No Weaver: Understanding Chinese Medicine*. New York: Congdon and Weed.
- Kövecses, Z. 1986. *Metaphors of Anger, Pride, and Love: A Lexical Approach to the Structure of Concepts*. Amsterdam: John Benjamins.
- Kövecses, Z. 2000. *Metaphor and Emotion: Language, Culture, and Body in Human Feeling*. Cambridge: Cambridge University Press.
- Krzyszowski, T. P. 1993. The axiological parameter in preconceptual image schemata. In *Conceptualizations and Mental Processing in Language*, R.A. Geiger & B. Rudzka-Ostyn (eds), 307–329. Berlin: Mouton de Gruyter.
- Lakoff, G. 1987. *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*. Chicago: University of Chicago Press.
- Lakoff, G. & Johnson, M. 1999. *Philosophy in the Flesh: The Embodied Mind and its Challenge to Western Thought*. New York: Basic Books.
- Langacker, R. W. 2000. A dynamic usage-based model. In *Usage-Based Models of Language*, M. Barlow & S. Kemmer (eds), 1–63. Stanford: CSLI.
- Maciocia, G. 1989. *The Foundations of Chinese Medicine*. Singapore: Churchill Livingstone.
- Mandler, J. M. 1992. How to build a baby II: Conceptual primitives. *Psychological Review* 99: 587–604.
- Manser, M. H. 1999. *Concise English–Chinese Chinese–English Dictionary*. Oxford: Oxford University Press.
- Manjali, F. D. 1997. Force Dynamics in language and cognition. www.chass.utoronto.ca/epc/srb/cyber/man4.html.
- Maturana, H. R. & Varela, F. J. 1980. *Autopoiesis and Cognition: The Realization of the Living*. Dordrecht: D. Reidel.
- Meng-Tsee-Khang, 2001. *Fundamentals of Tibetan Medicine*. Dharamsala: Mee-Tsee-Khang Institute.
- Milner, A. D. & Goodale, M. A. 1995. *The Visual Brain in Action*. Oxford: Oxford University Press.
- Porkert, M. & Ulman, C. 1982. *Chinese Medicine*. Cambridge, MA: MIT Press.
- Schipper, K. 1993. *The Taoist Body*. Berkeley: University of California Press.
- Sivin, N. 1987. *Traditional Medicine in Contemporary China: A Partial Translation of Revised Outline of Chinese Medicine (1972) with an introductory study on Change in Present Day and Early Medicine*. Ann Arbor: Center for Chinese Studies, University of Michigan.
- Talmy, L. 2000. *Towards a Cognitive Semantics: Volume 1: Concept Structuring Systems*. Cambridge, MA: MIT Press.
- Talmy, L. No date. Recombinance in the evolution of language. www.tech.plym.ac.uk/socce/evolang6/talmy.doc.
- Unschuld, P. U. 1985. *Medicine in China: A History of Ideas*. Berkeley: University of California Press.
- Unschuld, P. U. T. 1998. *Forgotten Traditions of Ancient Chinese Medicine: The I-hsüeh Yüan Liu Lun of 1757 by Hsü Ta-Ch'un*. Brooklin, MA: Paradigm.
- Unschuld, P. U. 2003. *Huang Di Nei Jing Su Wen: Nature, Knowledge, Imagery in an Ancient Chinese Medical Text*. Berkeley: University of California Press.
- Wiseman, N. & Ye, F. 2000. *A Practical Dictionary of Chinese Medicine*. Brooklin, MA: Paradigm Publishers.
- Wujastyk, D. 2001. *The Roots of Ayurveda*. New Delhi: Penguin Books.
- Ziemke, T. 2003. What is that thing called embodiment? Paper presented at the Annual meeting of the Cognitive Science Society, Sweden.
- Zlatev, J. 2007. Embodiment, language and mimesis. In *Body, Language and Mind: Volume 1: Embodiment*, T. Ziemke, J. Zlatev & R. Frank (eds), 297–337. Berlin: Mouton de Gruyter.

Get and the grasp schema

A new approach to conceptual modelling in image schema semantics

Paul Chilton

Get is a verb with a wide range of uses that are not obviously related to one another and has begun to attract attention from linguists in a broadly generative or formal semantics tradition (Haegemann 1985; Gronemeyer 1999, 2001; Fleisher 2006; McIntyre 2005; Manna 2004). Within Cognitive Linguistics there is as yet no comprehensive attempt to understand the meaning of this surprisingly complex verb (but see Hollmann 2003; Lee 2005; Bonnefille 2006). The present paper is an attempt to make progress towards a unified account, but the reader should be forewarned that it also proposes a novel theoretical framework. This framework relies heavily on spatial concepts formalised in informal geometric terms; crucially, the framework integrates foregrounding/backgrounding in discourse, temporal viewpoint, modal distance and directionality (see Chilton 2005, 2007 for previous expositions of the model). The fundamental principles of the approach are cognitive: it is proposed that the construction meanings associated with *get* are a conceptual category revolving around a prototype whose meaning is embodied in an image schema.

1. The analytical framework

In English the following sentence has become a common way of making a request:

- (1) Can I get a beer?

Without context, some speakers might interpret the subject here as occupying an agent role – that is, the speaker is proposing a movement that will result in the obtaining of a drink. Other speakers appear to understand the subject as occupying a beneficiary or receiver role. Clearly this ambiguity could be modelled in terms of predicate–argument structure and thematic role variations. However, there are many other meanings and constructions in which *get* is involved and many of these cannot be modelled in those terms. Assuming that a unified treatment is desirable, we need to look to a different theoretical framework. In this regard, (1) suggests something important about *get*. Its meaning can evoke the full “obtaining” cycle, in which the grammatical subject is agent, with the focus on the forward away-from-the-body reaching movement. Alternatively, it

can evoke a focus on the toward-the-body receiving part of obtaining, the end of which is the having of an object.¹

The puzzle that is the focus of this article is the following. How does the English word *get* get to mean so many things? The most challenging part of this puzzle concerns the question how *get* in its physical denotation (e.g. “Harry got the book off the shelf”) is related to its abstract modal sense (e.g. “Harry has got to read it by tomorrow”)? The approach to this issue makes the assumption, now usual in CL, that such meanings are associated with constructions, which are in turn associated with one another in a structured fashion. I am adopting the position that lexical items are meaningful because they are linked to cognitive frames (or domains) (Fillmore 1982; Fillmore 1985; Langacker 1987, 1991). Further, I am assuming that some frames are linked primarily with image schemas (Johnson 1987; Lakoff and Johnson 1999; Feldman 2006), for example prepositions *in*, *over*, etc. and some go beyond image schemas (e.g. *shoes*, *ships*, *sealing-wax*, *cabbages*, *kings* ...). It is an important property of image schemas that they can be transformed, including backgrounding and foregrounding or their components. Approximately the same idea has been around in psychology since Piaget, who postulated schemas arising in various forms in the different stages of infant development (cf. Piaget and Inhelder 1971). For Piaget, schemas arose from sensory–motor activity and were a kind of cognitive abstraction that could almost be described as concepts because they subsume particular instances. Grasping – a schema which is important for understanding *get* – is precisely one such schema studied by Piaget (e.g. 1952). In more recent work Mandler (2004) has adopted the Lakoff and Johnson notion of image schemas within a neo-Piagetian framework, studying the emergence of schemas in child development.

In the present paper I will argue that we can account for the semantics of *get* within what is essentially an image-schema framework. There are two key initial claims. The first is that a particular image schema needs to be invoked that has hitherto been relatively neglected in the CL literature – a schema related to grasping.² And the second is that a geometric formalism is a natural way to model the image schema itself and transformations on it that relate the different meanings of *get*.

The hypothesis is that the meanings of *get* are concepts that are embodied, i.e. meaningful a priori because they are in some fashion linked, developmentally or innately, to cognitive structures that are neurally instantiated in the human brain, including parts of the brain serving motor functions, perception and proprioception. There is an assumption in CL that such cognitive structures are in some sense “basic” and that they provide the ground for concepts that are either more abstract or more complex. Until recently, the linguistics literature on image schemas (Johnson 1987; Lakoff and Johnson 1999; Clausner and Croft 1999; Croft and Cruse 2004; and others who have followed them) has not had

-
1. There is another possible analysis, namely: (1) actually expresses the active cycle of obtaining and possessing (e.g. by purchasing) and that *get*, together with interrogative and modal, is a politeness formula.
 2. Lakoff and Johnson (1999) do of course discuss the metaphor UNDERSTANDING IS GRASPING and this is indeed a metaphor that reappears in one of the idiomatic meanings of *get* (“get it?”).

much to say about GRASP.³ However, cognitive linguists have now become interested in the findings of neuroscience, particularly the findings of Gallese, Rizzolatti and their associates (see Rizzolatti, Fogassi and Gallese 2000). These findings are independent of linguistics, but they provide a priori reasons for considering that the action schema of grasping is important for natural language semantics. Feldman (2006: 166–168) in outlining a theory of “embodied” meaning mentions, albeit rather briefly, that the meaning of the English word *grasp* is just the sensory–motor action schema for grasping, which is instantiated in the motor context. Gallese and Lakoff (2005) had already proposed a detailed image schema for the concept GRASP. The fact that there is some neuro-imaging evidence, that the human motor cortex is excited by the observation of grasping in others as well as in the execution of *grasping* by the self, suggests the existence of a schema that is independent of action. It is also known from neuroscientific work on the human cortex that specific neural pathways (known as the “dorsal” stream) control the coordination of vision for actions of the grasping type as distinct from vision pathways (“ventral” stream serving the identification of objects (Goodale and Milner 1996). Gallese’s theory of mirror neurons claims that grasping and handling neurons fire not only for actually performed actions but also for *observed* actions of this type (Gallese e.g. 2003). Gallese strongly suggests that something akin to a schema, in the sense the term is used in CL, is activated independently of actual performance.

2. Modelling the prehension schema: Reach-grasp-receive-have

The most specific hypothesis adopted in this article is that the meanings of *get* can be given a unified explanation if we postulate a dynamic cognitive image schema corresponding in part to the grasp schema proposed in Gallese and Lakoff (2005: 467). In more detail, prehension can be thought of as controlled by a composite schema involving a sequence of actions over time:

- a. location of an agent and of an object;
- b. movement of agent’s limb(s) (prototypically arm(s)) in the direction of an object;
- c. this action culminates in the application of force on the object in the form of some form of grasping;
- d. retraction of effector to some region of the self’s body such as mouth or trunk.

Let us assume that prototypically the effector limb moves to the edge of peripersonal space, and is retracted to the core of peripersonal space (on which see below). The hypothesis is that the meaning of *get* in its most basic spatio-physical (and earliest) sense is closely related to some such schema. Notice, however, that *get* does not in fact specify the exact nature of what I have referred to as the grasping action itself. The semantic schema for *get* is not the same as the global motor prehension schema. What I am proposing is that *get* schematizes the reach and retraction phases of a motor sequence in which an



3. Mandler’s (2004) list is: ABOVE-BELOW, AGENCY, ANIMATE MOTION, BEGINNING-OF-PATH, CAUSED MOTION, CAUSED-TO-MOVE-INANIMATE, END-OF-PATH, INANIMATE MOTION, LINK, PATH, SELF-MOTION, SOURCE-PATH-GOAL, SUPPORT, UP-DOWN.

agent moves an object into peripersonal space and implies the static possession phase. When the object is in peripersonal space, the agent can be said to *have* or *possess* it. The schematic properties of the cycle are egocentric: the personal space is presupposed; the first phase of the sequence involves application of force away from the subject and towards the object; the second phase involves the displacement of the object away from the object's first location to a new location, namely the subject. The schema thus entails a causal chain, the intentional exertion of force and a consequent movement, or path schema.

The kind of motor activity in question is sometimes described by cognitive neuroscientists in terms of coordinate systems located at effector limb, trunk, or retina. Indeed, the grasp schema can be represented by a configuration of vectors in a three-dimensional physical space. The length of reach is variable and combined with a variable direction – this is why it is appropriate to describe the prehension cycle in terms of vectors within a coordinate system. To summarise, there are two principal phases, which are motions of the arm, with different directions, one away from self and the other towards self. The reach–retract cycle takes place over time $\langle t_i, t_j \rangle$, in sequential scanning. Finally, there is an outcome, a state, in which the grasper possesses the object over time in their space. This state, like other states, can be represented as the unbounded set of vectors whose time coordinates are later than t_j (cf. Chilton 2007).

The “length of reach” can perhaps be motivated by another notion from neuroscience, that of peripersonal space (Rizzolatti et al. 1997; Holmes and Spence 2004). Personal space is the physical space occupied by the body, peripersonal space the space surrounding the body basically corresponding to the arm's reaching length. There are several relevant features for the discussion of *get* that are worth noting informally. There are obviously two ways an organism can get an object into its peripersonal space: it can either forcefully bring it into that space or move towards the object. Moreover, it has been observed by many researchers (see e.g. Holmes and Spence 2004) that the brain's representation of the extent of the peripersonal space may be variably increased by contextual factors. If one is looking for cognitive-linguistic parallels, the notion of peripersonal space may provide a concretisation of Langacker's “dominion” (1991: 176–180). More abstractly, I suggest that length of reach in peripersonal space provides motivation for the unspecified unit vector that is used in the diagrams below.

The overarching speculation of this paper is that the meanings of *get* are linked with the image schema depicted formally in the spatial geometric terms outlined above. Given that image schemas play a crucial role in linguistic meaning, the question of how to meta-represent the schemas and their components is an important one for CL. Diagrams are ubiquitous in the CL literature but while they are illuminating they are not obviously principled. Those of Langacker may constitute a consistent iconic system, but if so this system has never been axiomatised or made completely explicit. Almost all cognitive linguists (Lakoff, Talmy, Langacker) incorporate elements that correspond to formal geometries, topology and vectors in particular, frequently alluding to such features as boundaries, inclusion, directedness and distance. The present approach takes the position that if spatial relations are so fundamental in accounting for linguistically encoded concepts, an appropriate metalanguage for the description of these concepts can be provided by simple geometric formalisms, of which the most relevant for present purposes are vectors in co-

Figure 1a. Basic *get* schemaFigure 1b. Basic *get* and *have* schema

ordinate systems (cf. O’Keefe and Nadel 1978; O’Keefe 1996, 2003; O’Keefe and Burgess 1996; Zwartz 1997, 2003; Zwartz and Winter 2000). This idea is developed within what I call Discourse Space Theory (DST) (Chilton 2005, 2007), which develops an abstract three-dimensional (discursive distance, time, modality) *discourse space* (i.e. not a three-dimensional Euclidean space).⁴

One of the limitations of CL image schema theory is its static pictorial nature. As will be seen, what is needed, in order to model the various meanings of *get*, is a dynamic representation of the different directionalities of the reaching-retracting-having cycle. Informal geometric vectors provide a natural and simple method. I am not attempting a rigorous mathematical demonstration here, but I shall use vectors and coordinates very informally and put forward the proposal that, conceived in a certain way, they provide an illuminating means of analysis. Putting together the earlier remarks – that image schemas are simplified abstractions with respect to particular motor activities and that vector notation can capture the spatial dynamism involved – I propose Figures 1a and 1b as the image schema diagrams for *GET*.

The reach phase is a vector because it is directed and has length (unit length of the arm in peripersonal space). The r_2 vector points in the opposite direction. The image schema is dynamic because it unfolds over time. The notation is close to standard vector notation. The component vectors are labelled r as a mnemonic for *reach* and *receive* (or *retract*). The first component (r_1) can be interpreted as a force vector: energy is directed towards an object; the second (r_2) can be thought of as a path vector, by which an object is moved towards the agent’s self. Both vectors are continuous in time, and in a sense r_1 “causes” r_2 .⁵

4. The use of geometric formalisation to encapsulate the grasp schema is commonplace in robotics. To ape human reaching-grasping-receiving, engineers have to start with mathematical models that can approximate to directional movements of the human arm, involving, essentially, articulated rods of various lengths, angle changes, and rotations in several planes. Vector geometry is the main way in which such physical movements are modelled and programmed into the robotic system.

5. In passing, note that German *bekommen* (‘get’) has both a causative element (*be-*) and a directional element denoting movement toward the agent (*kommen*, ‘come’): *bekommen* means something like ‘cause object to come towards agent’.

Let us consider the relationship between *get* and *have*. Semantically it is trivial that *if X gets Y, X has Y*; empirically it is telling that *have* and Germanic cognates are cognate with Latin *capio*, seize (or get). This relationship can be captured roughly as in Figure 1b, which depicts a fuller schema, including a phase in which an object is “possessed” or “had”. This too is a spatial relationship between object and possessor, but one that is the outcome phase of the action depicted in Figure 1a. Since it is a state, it is indicated here as an indefinite set of vectors relating the grasper-possessor to the object.⁶ It may seem counterintuitive that the arrows point from A (the person who is getting) and the object B. The reason for this is the mathematical vector notation, in which the arrows can be understood as position vectors “locating” B at A. In general I treat possession as position, notated in this way.

In the figures, the head-to-tail contact point between r_1 and r_2 is where the grasping occurs. The grasping action of the hand and fingers is in itself of course a schema in its own right but I have not attempted to include it here. While it is relevant to the semantics of the verb *grasp* it does not appear to be relevant to the semantics of the phenomena we are concerned with in this paper. The overall picture is simplified compared with the motor activity of real-world reaching-grasping-retracting (as are CONTAINER, PATH, etc.), but that is precisely the point of a schema (cf. Gallese and Lakoff 2005: 10–12 and Feldman 2006: 167 on parameterisation of complex motor schemas). It is also crucial to note that the schema is input to cognitive operations (abstractions, reductions, shifts, metaphors) that give rise to varying lexical meanings. Stating this slightly differently, the hypothesis is that the brain uses an image schema something like the one sketched out in Figure 1, abstracted from complex motor neural systems, and underlying the semantic coding that constitutes the meaning of *get*. As argued by Feldman (2006), image schemas are frequently elements in more complex frames.

The *get* schema is associated with a wide range of human activity frames that involve environmental and social knowledge. Thus, for instance, *John got a new car* involves a specific kind of commercial transaction frame, *John got a PhD* involves yet another frame, and so on.

The different meanings of *get* recognisable across particular contexts are assumed to be interrelated. To see these interrelations in terms of the basic image schema just outlined, we have to assume various operations on it, principally the weaker or stronger activation of particular components, as well as its extension beyond the basic physical meaning. The image schema proposed for reaching-grasping-retracting-having can be located within the three-dimensional discourse space postulated in DST. This is a crucial feature because it means that the modal dimension can be involved. It is in fact this feature that enables us to explain the cognitive-linguistic relationship between certain of the different abstract meanings of *get*.

6. Note that in DST the vertical axis is not space but “discourse distance” (d): some elements, e.g. expressions placed later rather than earlier in clauses, thus grammatical objects in English, are more distal or backgrounded. The object spatially moved under the semantics of *get* still has its discourse d-coordinate for the implied *have* schema.




Figure 2. Receiving phase of basic *get*

We are now in a position to explain (that is, propose explicit models for) the ambiguity of (1). Simply put, the interpretation of (1) in which the speaker is agentive and is asking if he or she can physically take hold of a beer foregrounds r_1 (cf. Figure 1). The reading in which the speaker is asking some other agent to transfer a beer to the speaker places the speaker in the role of recipient and foregrounds r_2 . The model is thus essentially directional; in fact it is deictic, since r_1 and r_2 involve movement from and towards the self, respectively. This might seem an odd way to model theta-roles such as agent and recipient. The point of the present approach is, however, to capture the spatial and the bodily basis of linguistically coded conceptualisation, including theta-roles. Whereas writing “agent”, for example, provides an arbitrary and explicit symbol, providing a geometric diagram brings us one step closer to spatial experience from which such roles, it may be claimed, are derived. The notation for the first (agentive) meaning prompted by *get* in (1) is as in Figure 1a; the notation for the reception meaning is as in Figure 2. Because the schema carries with it an implied *have* schema, it is equally possible to say “Can I have a beer, please”. In this case, the link with the prehension part of the image schema (r_1 and r_2) is backgrounded or weakly activated.

3. Basic prehension: *get*-NP-PP_{source}

Let us assume that there is a basic sense of *get* which corresponds directly with the abstract image schema of Figure 1. *Get* in this sense denotes an action “within reach” (i.e. within peripersonal space) in which the arm extends away from the agent’s body towards a target object, grasps the object with the hand, retracts the arm towards the body and subsequently *has* that object. I take sentence (2) to be an example of such a meaning:

- (2) John got an apple from the tree.
- (3) John got an apple from Eve

(3) can be read as implying that Eve gave John an apple – the case of r_2+h . John is not an agent in the getting, but a recipient.

In the following analyses I use diagrams that represent the meaning of utterances in abstract discourse space, the three-dimensional deictic space outlined in Chilton (2005).

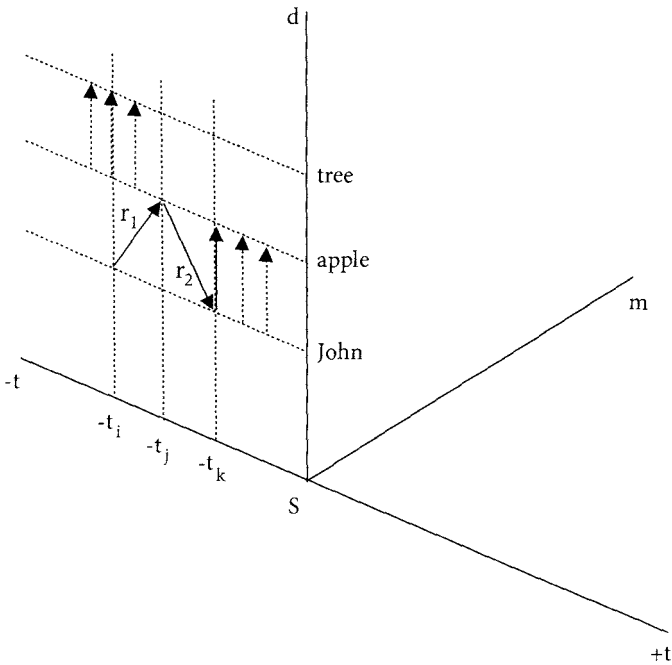


Figure 3. John got an apple from the tree

As in Figures 1 and 2 there are two axes, conceptual distance (d) from the speaker S and time (t) centred on the point of utterance. The third axis (m) represents modal distance, realis being close to S and irrealis distal to S . The uses of the m -axis for modelling *get* will become apparent later.

In Figure 3, we have the apple positioned at the tree coordinates at time $-t_i$ and at John coordinates at time $-t_k$. From S 's viewpoint, the *get* event is past and takes place entirely in the realis plane. The two locative states – prior location at the tree and subsequent location at John (i.e. possession by John) – are shown in Figure 1b as samples of indefinite sets. The word order of the sentence corresponds to the relative conceptual distance from S , the sentence focus (“discourse distance”). The length of the vectors (arrows) in these diagrams is an arbitrary unit that can be modified by pragmatic or lexical environment related to peripersonal space, as suggested earlier.⁷

A final but important point needs to be made about the full basic *get* schema. Intuitively, basic *get* presupposes the desire or intention of the agent to acquire an entity. This meaning is specifically tied to the r_1 vector. Sentence (2), “John got an apple from the tree”, assumes John’s intention and desire unless specifically overridden (e.g. by adding “under duress”), while “John got a medal” (cf. (6b) below) may or may not presuppose desire and (4) does not normally presuppose desire. These differences arise because of pragmatic

7. Figure 3 is essentially a more explicit representation of Pustejovsky’s (1988) formulation of event structure cited by Gronemeyer (2001:7).

knowledge including values (concerning medals and slaps) and are represented theoretically here if we say that r_1 is present in (2) and possibly present in (6b) below.

- (4) John got a slap in the face

In (4) only r_2 , i.e. the reception part of the grasp cycle, is present. Since r_1 is assumed to be a force vector, it appropriately stands for effort and by implication for desire and intentionality, while r_2 is constructed in the present theory as a path vector directed towards self. Finally, the schema does in fact seem to privilege or foreground the end state, of reception and possession.

4. Getting and having: *has got*-NP

The verb *get* includes a reference to the movement of an object from one location to another. The source location may or may not be mentioned in a PP. Because the image schema includes a *having* phase as an intrinsic part of prehension, two alternative expressions for possessing are available in English:

- (5) a. John has got an apple [a car, a degree, etc.]
 b. John has an apple [a car, a degree, etc.]⁸

(5a) can be read as denoting the whole prehension cycle, though (5b) cannot; (5a) can also be read without r_1 (John has been given an apple, for example). Further, it appears that the image schema in Figure 1 can background both r_1 and r_2 and foreground just the having condition. These meanings can be separated out by lexical or pragmatic environment, as in the following:

- (6) a. John got (himself) a beer from the fridge r_1 and r_2 and {h}
 b. John got a medal from the President r_2 and {h}
 c. John's got a degree in theology. r_2 and {h} or {h}

One reading of (6c) would be the present perfect form of “get a degree in theology”; (6c) differs from “John has a degree in theology” in that the latter is not open to ambiguity. The possibility of separating out these different phases of the image schema is of considerable importance in understanding how the various meanings of *get* are conceptually related. The varying meanings arise from the sentence context, which operates on the basic *get* schema.

4.1 Beyond peripersonal space

In diagramming (2) we assumed that John was “within reach” of the tree, though of course other readings are possible. In examples (6a) and (6b), the schema is extended beyond the

8. For some speakers *has* and *has got* may not be alternatives if the NP is abstract: ?*John has got a thirst for life*, etc.

peripersonal reach-space that I take to be the limiting space of the fundamental concept. In (6a), John might have to walk to the fridge before it is in his peripersonal reach space; in (6b) physical reach-receive space is not relevant; in (6c) no spatial relationships are involved at all. All of these variant interpretations result from the cognitive frames associated with particular lexical items. Extension beyond peripersonal space may in the first place be naturally metonymic: self-movement may be required for the purpose of prehension. Beyond this, getting can involve complex processes, including social processes of various kinds, by means of which the agent ends up *having* some object.

4.2 Beneficiaries of *get*: $\text{get-NP}_{\text{ben}}\text{-NP}(-\text{PP}_{\text{source}})$ and $\text{NP-get-NP-PP}_{\text{ben}}$

Ditransitive examples may be extensions beyond personal space, but not necessarily so. McIntyre (2005:4) describes ditransitive *get* examples as “augmentations” of the VP with a beneficiary. However, this assumes that examples like (7a) and (7b) are completely synonymous. In CL, I take it that we would want to capture some difference in conceptual construal.

- (7) a. John got/fetched Mary a book
 b. John got/fetched a book for Mary
 c. John got/?fetched a book to Mary

There are three possible interpretations for both (7a) and (7b). In one reading John gets the book and gives it to Mary in one continuous action with two consecutive phases, getting and giving. In another reading, John gets the book because he wants to give it to Mary at some point in the future, relative to utterance time, as pointed out in other terms by McIntyre (2005:404) and in construction-grammar terms by Goldberg (1995:32) and Croft (2000:4). In a third reading, Mary has asked him to get the book and give it to her, and the getting and giving may or may not be continuously consecutive. In all three of (7a), (7b) and (7c), *fetch* appears restricted to an event in which the recipient comes into immediate possession of the object. Thus in (7c) there is a possible, and perhaps most probable, interpretation arising from the preposition *to*, that Mary is distant from John, that John is not necessarily the agent of the sending to Mary of the book, and even that John is not the physical getter of the book in the first place.

It would take us too far afield to pursue the appropriate models for all these cases in detail. However, there is a general point to make. The meaning of the lexical item cannot readily be specified in terms of a combination of componential primitives but rather requires the holistic notion of schemas. Furthermore, it is necessary to invoke a modal dimension, since in the *get* readings of (7a) and (7b) outlined above, the constructions do not entail that John in reality gave the book to Mary. Rather than treat this possibility as semantic entailment, however, the model regards the non-reality as a mental representation on the part of the utterer – or on the part of a hearer interpreting the utterer’s intended meaning. Figures 4a and 4b include the modal axis of the discourse space and show two meanings for the *get* case in (7a).

The diagrams need some commentary. Figure 4a models an interpretation of (7a), and in general of the construction NP1-get-NP2-NP3 that corresponds to an event in which

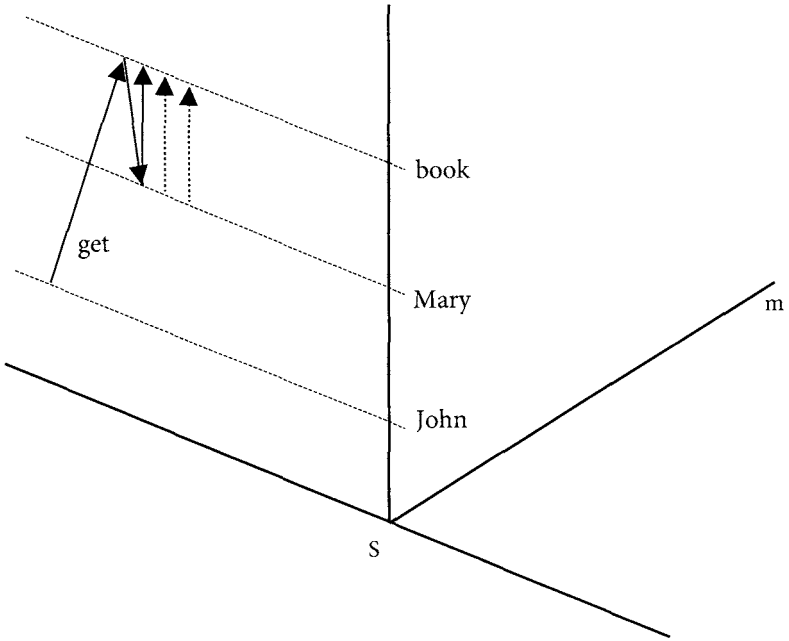


Figure 4a. (7a) John got Mary a book




Figure 4b. (7a) John got Mary a book

John gets the book and continuously transfers it to Mary. To model this we have to introduce another transformation of the basic image schema – one in which r_2 translates the object not to the starting coordinate of r_1 , i.e. to the agent, but to some distal entity, in this case Mary. Discourse entities are indicated by coordinates on the d-axis, which represents relative discourse distance (often reflected iconically in word order).⁹ Mary subsequently *has* the book – here depicted by three sample vectors of an indefinite set.

Figure 4b corresponds to an event in which the speaker S can truthfully assert that John obtained possession of the book, but cannot truthfully assert that he transferred it to Mary. Moreover, *for* seems to imply something about John's desire or intention as conceived by the speaker S, not about S's own knowledge of the facts. The way DST models this may look complicated but it captures detail that takes up many words in other accounts. From the second reading of (7a), and also from that of (7b), we know that S thinks John has the desire or intention to transfer the book to Mary at some point t_k later than t_j ; the diagram gives a precise position for the *give* phase for the sake of concreteness. The crucial point here is that in asserting (7a) in the required sense S is making an assertion about some state of the world as conceived by John, a state of the world which, from S's point of view is, at utterance time, neither true nor false. From S's point of view, John might give the book to Mary or he might not. So, the diagram includes an embedded set of axes (a world, if you will) whose origin is coordinated on John. The mid-point of the m-axis is epistemically neutral, i.e. neither true nor false from S's viewpoint.¹⁰

The ditransitive *get* construction thus requires at least two discourse space models that include a modal dimension. The same approach can probably be generalised to other subcategories of verbs that occur in English ditransitive constructions.

5. *get*-NP-PP_{goal}

Because they involve PPs, sentences such as (8) appear syntactically similar to (2).

- (8) a. John got the book to Mary
 b. John got his suitcase on the truck
 c. John got dust in the camera

There are, however, several differences conceptually, apart from the distinction between source and goal. A key difference is that the basic prehension schema illustrated in (2) is restricted to the preposition *from*, *from* under, etc. and *out of*, while the PP_{goal} constructions in (8) admit a bigger set of spatial prepositions. Additionally, the prior location may or may not be specified in the sentence (cf. *John got the parcel from Manchester to London*)

9. This is how dative/particle shift can be handled in a cognitive account, i.e. the construal difference between e.g. *John gave Mary a book* and *John gave a book to Mary* can be captured. There is no space to pursue the details here.

10. Note that the configuration of the vectors for *give* resembles that for *get*, as we should expect from the semantic overlap between *getting somebody something* and *giving somebody something*, but I do not pursue this point further here.

or in surrounding discourse, or there may be a default that the source location is in the space of the subject. A further difference is that examples similar to (2) involve physical contact between agent and object, whereas examples of the (8) type, including (8b), do not necessarily entail physical contact between agent and object. (8b) is ambiguous in three ways. First, John exerts some force and this force is direct (e.g. he heaves the suitcase onto the goal location). Second, John exerts some force indirectly on the suitcase (e.g. he persuades or pays somebody to heave it on for him). In this latter reading John does not handle the suitcase himself, although the basic schema, instantiated in (2), is defined in terms of physical interaction between body and environment. Like (2), type (8) does involve the physical motion of an entity to a location. I will assume that the r_1 vector in subsequent diagrams can be taken schematically to include indirect force and that the model for sentences like (8) can be derived from the basic mage schema of Figure 1, by means of a simple geometrical transformation that relocates the end point of r_2 at some point with a d-coordinate other than that of the agent. Thirdly, (8c) has an unintentional-only reading – something like “John was lucky enough to get his suitcase on the truck”. In this instance the force phase (r_1) appears to be absent, since John is taken not to have consciously initiated the event, even though there is some residual sense that John is responsible, or even the (unconscious) cause. It seems clear that the model should represent a “reception” or “arrival” phase r_2 . We will consider these issues further below.

Figure 5 is an attempt to model the semantics of sentences like (8a) and (8b) but not (8c). What these sentences have in common is a goal-location that is distal relative to the NP subject (*John*). The configuration of vectors is a derivative of the base schema (Figure 1a) and is related to Figure 3a. What is important is that Figure 5 incorporates

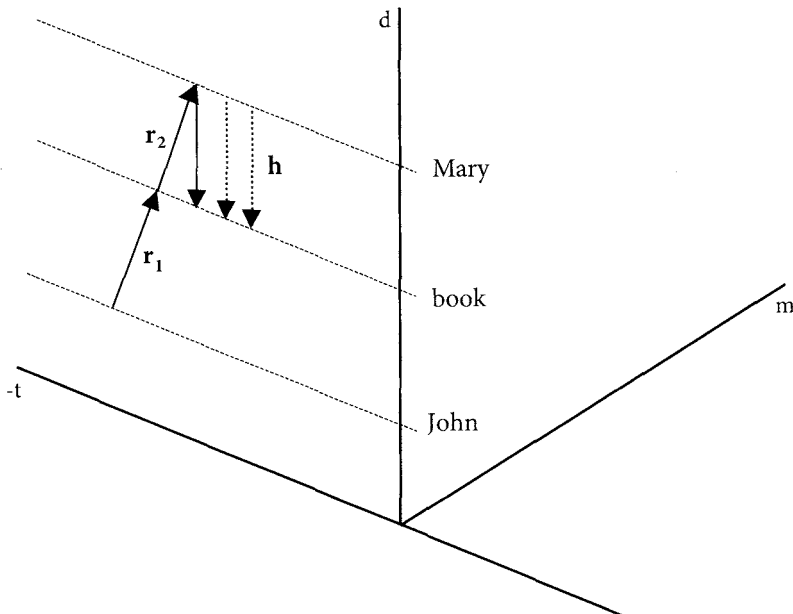


Figure 5. (8a) John got the book to Mary

both r_1 and r_2 . That is to say, it includes a representation of voluntary intentional force-application and a phase of completion r_2 , where the latter arises from the “reception” phase of the base schema. Leaving r_1 out of the model would give a representation of *Mary got the book*. The set {h} is *Mary has (got) the book*.

In this analysis, *contra* McIntyre (2005: 404–406), it is implicit that uses of *get* which imply hindrance or difficulty experienced during the action are actually agentive. McIntyre makes the claim that “hindrance-*get* uses” are non-agentive on the basis of rather weak evidence, mainly the *almost* test, and by using the terms *agentive*, *intentional*, *causative*, *responsibility*. These latter terms, while to some extent standard in semantics, do in fact stand need of conceptual clarification, something that cannot be attempted here.

In *get* sentences the scope of *almost* is sometimes narrow, sometimes wide:

- | | | | |
|-----|----|------------------------------|------------------|
| (9) | a. | I almost got the fire out | [narrow only] |
| | b. | I almost put the fire out | [wide or narrow] |
| | c. | I almost got the car out | [wide or narrow] |
| | d. | I almost got the car through | [narrow] |

My analysis of these differences is as follows. A “hindrance” reading is triggered either by the preposition in the PP (e.g. *through*) or by the frame evoked by the lexical material or by some context. Putting out fires implies difficulty and (9a) denies completion of the action but cannot have the wide reading that denies I even started the act of putting the fire out, in contrast to other verbs, e.g. (9b). In (9c), the reading is wide if the contextual frame concerns travelling, narrow if the car is in a ditch. (9d) suggests hindrance because of the preposition. Now, McIntyre’s (2005: 405) claim is that the *almost* test demonstrates that “hindrance-*get* VPs do not denote causing events or agentive acts”. The logic seems to be that verbs like *put out* in (9b) denote accomplishment type *Aktionsarten* and the *almost* construction containing *put out* can be interpreted as meaning: either (i) the subject formed an intention of putting out the fire or prepared at least mentally to do so but then did not perform the action, or (ii) the subject began putting out the fire but did not succeed in causing the fire to go out. In (9a), the first of these meanings is not available. McIntyre’s argument seems to be that this fact indicates that *get*, specifically in a “hindrance” context, does not have agentivity or causativity in its semantic structure. Presumably, the idea is that since wide-scope *almost* denies agentivity and causation of an outcome, the verb *get* does not have any agentive component to deny. There are two problems with this. One concerns what is meant by agentive: one can act (be an agent) without causing an outcome. Second, the effect of *almost* need not be on agentivity. Indeed, what is at issue here is the semantics of the *almost* construction itself. To put it summarily, my claim would be that in *get* constructions *almost* actually denies not that the subject acted but denotes rather a late phase of the event denoted by the verb, whether an accomplishment or an achievement verb.¹¹

In terms of the framework we have outlined for *get*, the effect of the *almost* construction is to suppress r_2 , the completion phase of *get*. In general, in constructions (not in-

11. As is perhaps indicated by the etymology of *almost*: most of all an action is foregrounded by this word.

volving negation operators like *almost*) that use lexical items whose frames involve background knowledge about attendant difficulty, etc., there will be r_1 , which we have defined as a force vector and interpreted as indicating willed intended action by an agent. Thus the reason why decontextualised (10) has a “hindrance” reading is that the reader has to construct an interpretation of *across the room* which accounts for the agentive-initiative feature (i.e. the r_1) of *get*.

(10) I got the book across the room

This is surely what the decontextualised display sentence (10) tells us. It is not that we first of all have, say, a crowded-room context, and then we know we can correctly use hindrance-*get*. Rather, the hearer infers from hearing (10), perhaps in a conversation the next day, that the room was crowded and makes such an inference because *get* inherently contains the initiating agentive meaning. In the present analysis I am assuming that the full representation of *get* involves the force vector r_1 and so in (10), it is the presence of r_1 in the semantic specification of *get*, clashing with the normal expectation of crossing rooms with books, that forces a construal of *the room* as meaning, say, a room full of people or that the book was enormously heavy. Note the PP *across the room* does not in itself imply hindrance; rather, to repeat, the r_1 in *get* makes the reader fill out the meaning of *room* (or *book*). Some contexts are inherently “hindrance contexts” and are compatible with *get*; some are not and are made into “hindrance contexts” by the presence of the r_1 in the structure of *get*.¹²

To sum up, the “hindrance” reading is a function of lexical or non-linguistic context. The *almost* test does not tell us that there are two types of *get*, one with a hindrance reading, the other one causative-intentional. Rather, this “test” tells us that *get* has a schematic meaning that interacts with constructions. This can be seen from (9c), which is ambiguous with respect to the presence or otherwise of a hindrance interpretation. The latter is not a test but a construction with *get* in it. The question is how to represent the effect of *almost*. Sentences like (9c) let us infer that the action as a whole did not take place and that the subject formed an unrealised intention to bring about the event. Alternatively, we can infer that the subject started the getting out action but did not complete it. DST can model both these situations by using the m-axes. This is not diagrammed here, but can be summarised as follows: in the narrow reading, r_2 would appear in the irrealis or counterfactual plane

12. There is an important ancillary issue connected to these characteristics of *get*, namely the effect on particle order, discussed in depth in McIntyre (2005):

- a. I got the rusty key in
- b. *I got in the rusty key
- c. I got out my wallet (It was in my pocket)
- d. I got my wallet out (It was in my pocket)
- e. I got my wallet out (I had dropped it down the drain)
- f. *I got out my wallet (I had dropped it down the drain)

Space prevents discussion of this matter but a cognitive account might pursue the following points. Phrasal verb *get out* rests on the CONTAINER schema and having access or control to something. Where *get out* expresses effortful extraction from a containing location, *out* occurs at the end of the process of extraction and iconically at the end of the clause.

located at the extreme end of m ; in the wide reading *almost* would appear as an operator moving r_1+r_2 to the irrealis plane, over a time interval.

Let us return now to (8c), repeated as (11):

(11) John got dust in the camera

Unless John is malicious, (11) is construed as an unintentional act (*John got dust in the camera but he didn't intend to*). How can the meanings of (11) be modelled in the present framework? In a first approximation, we could envisage simply using Figure 5, minus r_1 . The camera, distal to John, receives dust; John does not initiate an act that causes this result, so there is no r_1 . However, the resulting diagram seems more appropriate for:

(12) Dust got in the camera

where dust is discursively closer, camera discursively distal. This is tantamount to a movement meaning of *get* – a meaning that will be discussed in detail in a later section. (11) requires a different analysis, not only for this reason, but also to try to capture the meanings of (11), which include the intuition that John is in some sense responsible for the dust being in the camera, though he did not cause it. Moreover, the syntax of (11) is not obvious, since neither *John [got[dust in the camera]]* nor *John [got dust[in the camera]]* seems appropriate. The syntactic relationship between *dust* and *in the camera* is what is at issue and here I propose a conceptual resolution of the matter: *John* and *the camera* are in some sense, conceivably in an abstract quasi-spatial sense, co-located, at least in the sense that the camera is in some sub-region of John's (maybe extended) peripersonal space (cf. discussion in Section 2 above) or, to use Langacker's term, in his "dominion" (Langacker 1991: 176–180).¹³ The irresolvability of the syntax may be a reflection of this conceptual part of the construction's meaning.

These notions are incorporated in Figure 6, where John is a recipient, as in the prototype image schema, though he does not initiate an action: there is therefore no r_1 . Co-location of *John* and *the camera* is captured by their having the same coordinates on the d -axis. This is justified by, and in a rather sparse way relates to, the intuition that the subject's responsibility for the resultant state is part of the construction's meaning. For if the camera is in John's peripersonal space, and in human contexts one's space is the space over which one has "dominion" or control, then John may be regarded as responsible for it. These conceptual possibilities, we might hypothesise, have to do with rich cognitive frames linked to human bodies, the space around them and social knowledge.

In similar sentences syntactic co-indexation is required, as noted by McIntyre (2005: 408):

- (13) a. The camera_i got dust in it_i
 b. John_i got dry rot in his_i house
 c. *John got dry rot in Bill's house
 d. John got a nail in my tyre

13. Langacker's definition (1991:547) is highly suggestive: "The set of entities (or region comprising them) that a particular reference point allows one to make contact with."

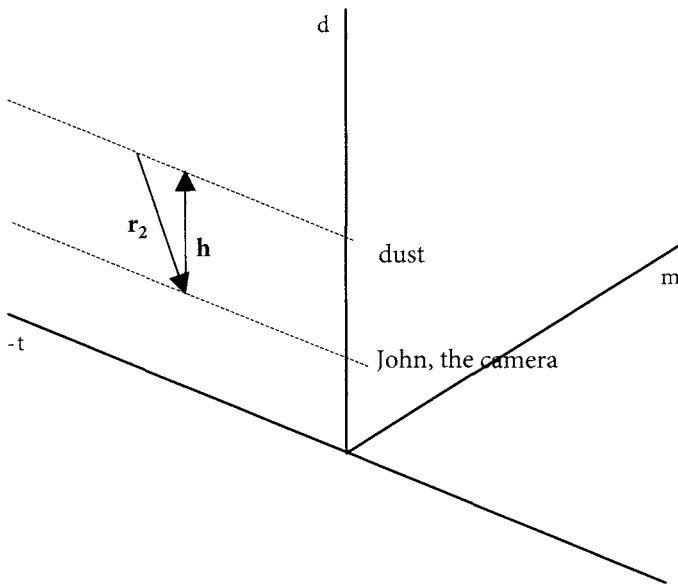


Figure 6. John got dust in the camera

(13a) can be modelled by Figure 6, without the co-location of John.¹⁴ The reason for the unacceptability of (13c) certainly arises from the associated cognitive frames. Even if Bill were not living in Bill's house, what we assume about dry rot makes it difficult to ascribe responsibility to John. However, (13d) has no co-indexation and is conceptually acceptable for a context in which John is driving my car. Nevertheless, (13b, c and d) strongly suggest a concept of responsibility, for a peripersonal region is relevant here: this would explain (13c), while (13d) is explained if John has temporary dominion over my car while driving it, surely the default interpretation.¹⁵

It may seem strange that the concept of responsibility should need investigation at all in order to account for purely linguistic explanation. However, many descriptions of certain *get*-constructions invoke it, so it is worth asking why it should become involved in uses of this verb. The notion of a quasi-spatial domain of personal responsibility has already been sketched: responsibility seems to involve spatial conceptualisation, specifically conceptualisation of peripersonal space. Responsibility for some event conceptually involves both agency and intentionality. For a complete answer we would need to investigate the phenomenology of intention and responsibility and their relationship – which we cannot do here. However, we may note that responsibility for an event can be attributed without intentionality, but not vice versa. In terms of our modelling of *get*, we can say that

14. Space forbids a consideration of why *the camera got dust* is unacceptable in contrast to (13a). Animacy may play a role: *the dog got fleas*, **the table got a stain*, but note *the house got a lick of paint*.

15. One inconvenient aspect of the DST modelling should be acknowledged: the principle that d-axis reflects discourse order/distance in word order is overridden by the cognitive spatial co-location. Reflexives raise the same problem.

if we interpret the force vector r_1 as entailing intentionality, those meanings of *get* represented as $r_1 + r_2$ entail the responsibility of the denoted agent for an event. Those constructions represented by r_2 alone might be expected to have non-responsibility readings. Why then do we still get the residual notion of responsibility with r_2 alone, as for example in (11)? The answer to this might, in addition to the presence of a domain of responsibility, be the vestigial presence of intentionality. The r_2 constructions are derived from the full $r_1 + r_2$ form and thus remain somehow cognitively linked and still weakly active.

These are speculative matters, but on a more technical point, one theoretical implication of the formalism used here is the following: there can only be one endpoint, or in the present terms “receiving location”, for r_2 . That is, the recipient in (11) cannot be both John *and* the camera where John and the camera have separate coordinates on the d-axis. This means in effect that the model predicts that *get* constructions can have no more than one conceptual endpoint. This principle should not be understood in terms of the presence of PP_{loc} in sentences, but in terms of the conceptualised meaning of the sentence. Putting this in different terms, it should be anomalous to find conceptualisations like *X got Y Z_{loc}* unless *X* and *Z_{loc}* are conceptually co-located. And this is indeed what the above examples show to be the case, although what precisely is to be understood by “conceptually co-located” remains for further investigation, perhaps along the lines suggested in the above paragraphs.

6. Intransitive *get*: *get-Part_{loc}* and *get-PP_{loc}*

We have seen how the *get* schema can be transformed into a configuration in which an object is translated not to the self or the self’s proximal space but to a distal location (Figures 4a and 5). We now have to confront intransitive constructions, which are processed as denoting the movement not of the entity denoted by the object NP of *get* but of the entity denoted by the subject NP (cf. (12)).

- (14) a. John got in/out
 b. John got on/off the bus
 c. John got to London
 d. Dust got in the camera (cf. (12))

The particles and PPs in (14a–b) are primarily stationary-locative but interpreted here as directional; in (14c–d) the prepositions are semantically directional. (“Particles” here can be viewed as prepositions whose complements are provided by discourse referents already in the context.) In all four cases *get* is intuitively understood as motion of the entity denoted by the grammatical subject. That is, in this meaning *get* no longer denotes movement of a body part (arm) and/or translation of an object towards the agent, as in the prototype schema, but movement of the entire agent (or non-agentive entity in (14d)). What needs to be explained is how the basic prehension schema is reconceptualised in this fashion.

One option is that movement *get* arises metonymically from the basic image schema because of its extension beyond peripersonal reach-space: spatial relocation of the agent is part of getting and it is this part of that movement which *get* selects. In this case it ought

to be roughly synonymous with *go* but this is clearly not the case, since *John went on the bus* is not synonymous with *John got on the bus* (14b), and a similar point can be made for (14c) and (14d).

It therefore seems reasonable to consider another explanation and a different transformation of the basic image schema. This approach leads us towards seeing *get* as an aspect operator. In Aktionsart terms, *get* induces an achievement reading, in the case of a PP headed by *to* and other PPs. That is, *get* produces a cognitive focus on the completion phase of a process, as the *almost* construction shows:

- (15) a. John almost got to London [narrow scope only]
 b. John almost walked to London [wide or narrow scope]
 c. John almost went to London [wide scope only]

The inference in (15a) is that John had already started out to London. (15b) is ambiguous in having either wide or narrow scope reading. (15c) shows that *get* is not equivalent to *go* in that *go* has only wide scope. The PP in (15) encodes both the end location and directed movement. What is interesting is that *get* does not seem *per se* to encode movement. Rather, in (15a), movement is inferred by the reader on the basis of the PP, and *get* adds achievement action typing to that inferred meaning. That the semantics of the PP is significant is shown in (16):

- (16) a. John almost got on/off the bus [wide or narrow scope]
 b. John almost got up the stairs [narrow only]
 c. He almost got through the tunnel [narrow only]

In (16a) John considered leaving the bus, e.g. because it was so dirty; alternatively he tried to leave it but was unable, e.g. because there were people in his way. In (16b), however, a wide scope reading is less easy to accept and (16c) even less easy.

There is a further important meaning. In cases like (15a) there appears to be a “trying” presupposition. However, this is not a permanent feature of *get*, as (17) shows:

- (17) I will send you an email as soon as I get to/arrive at/reach my office.

Here *get to* seems to be near-synonymous with *arrive at*. “Trying” does not seem to be an automatic inference. What *get* does is indicate the termination of a goal. The only inference is that the subject moved from one location to another; *get* (like *arrive* or *reach*) enables a speaker to assert this without referring to the manner of travel.¹⁶

In general, certain frequently observed meanings associated with *get* (effort, movement and hindrance) appear to arise from the interaction of *get* with contextual elements, i.e. other conceptual frames. I do not propose to pursue further the effect of different prepositions, lexical material and contexts in (15) and (16), but merely try to outline how the present framework can represent the two meanings that typically arise in different cases. These are what we might call the arrival meaning (narrow scope with *almost*, cf.

16. As Slobin (1996), Talmy (1985), and now many other scholars, have observed, the English lexicon has a strong tendency to oblige speakers to refer to manner of motion in the main verb: *get* makes it possible for English speakers to avoid this.




Figure 7. John got to London

(14a–d)), and the agentive meaning (wide scope with *almost*, cf. (14a) and (14b)). We therefore need an account that is appropriately schematic, i.e. has an abstract conceptual structure that will yield particular conceptual effects in particular grammatical (and contextual) environments.

Figure 7 is like Figure 5, except that it has two not three discourse referents. Since intentional effort is implicit in the sentence, the force vector r_1 is present, and since the sentence certainly implies arrival, r_2 , derived from the reception vector of the base schema, is also present. To model *John tried to get to London*, the DST framework would show the arrival phase r_2 in the counterfactual plane, r_1 in the realis plane. The addition of r_1 and r_2 is the completed event structure for an achievement verb, as discussed in Chilton (2007). The r_2 can be punctual or, as here, happening over time: this does not affect the general argument made here. The entailment “John is at London” is shown by location $\{h\}$ corresponding to possession $\{h\}$ in Figure 2 for the base schema, a consistent result, since have-possession is understood as location.

Note that r_2 can be derived from r_2 in the base schema by geometrical reflection of r_2 about the horizontal. Accordingly, the direction of r_2 is reversed: John goes to an entity (London), not an entity to John, while a time course is preserved; the direction of vectors in the set $\{h\}$ is reversed in direction also, representing the positioning of John in relation to the entity London, not the entity at John (possessed by John). By this transformation of the base schema, instead of John as recipient we have London as “recipient”. This wording is not inappropriate, since recipient is also a kind of goal, transfer of objects being itself a special case of the source-path-goal schema. Note that r_2 needs to be semantically filled out for the preposition *to*. Figure 7 also stands, *mutatis mutandis*, for (16b) and (16c), with appropriate changes to the second discourse referent (*stairs, tunnel*) and for the different




Figure 8. John almost got on the bus [narrow scope reading]

spatial relations expressed by *up* and by *through*. It is important to note that DST is not intended to represent the entire semantics, only the fundamental discourse space. A further advantage – or motivation – of this way of modelling the construction in question is that it allows for the representation of the reflexive variant of (15c): *John got himself to London*. To capture this variant, Figure 7 would include a referent on the d-axis for *himself*.¹⁷

With regard to sentences (14a) and (14b), which can have wide or narrow scope readings with *almost*, they are treated here as both having r_1 elements, i.e. effort and intention. However, in constructions where it appears, the presence of *almost* does make a difference. Consider (16a). If the inference is that John began to get on but did not complete the action, the modelling is as in Figure 8. If the effect of *almost* in (16a) is to deny that John did not act but that he formed an intention to act, then both r_1 and r_2 are in the irrealis plane under the effect of an *almost*-operator.

17. There remains the question of how DST deals with reflexives in general – a matter for future clarification; cf. note 15.

7. State as complement: *get*-Adj/V_{past part} and *get*-NP-Adj/V_{past part}

States expressed by adjectives and by past participles overlap morphologically and semantically in English, a fact used by Fleisher (2006: 230) and Gronemeyer (1999: 6) in their studies of *get*. Intuitively, the meaning of past participles includes an event resulting in the state and occurs in the constructions referred to as “passive *get*”. Fleisher, following Gronemeyer argues that *get* with participial complement develops diachronically from *get* with adjectival complement. The aim here is not to account for historical change but to understand the conceptual ground for the relationship.

Both adjectival and participial constructions can be incorporated in the present framework, provided the step is taken of treating properties as states that can be given coordinates as referents in the discourse space. Further, this spatialisation of state concepts is compatible with observations familiar in cognitive linguistics. For example, movement verbs frequently appear cross-linguistically and diachronically as change of state verbs. Such uses may involve NPs: *she went from poverty to riches, from a state of ignorance he arrived at understanding*, etc. They may involve non-finite clauses: *she came to believe his story*, etc. And they may involve adjectives: *the sky went from grey to blue, his condition went from bad to worse, the patient went from depressed to elated*.¹⁸ An ancillary though not conclusive observation is that attributive adjective constructions with *be* (e.g. *Jane is courageous*) are sometimes close in meaning to locational constructions with *have* and a morphologically related noun (e.g. *Jane has courage, Jane has intelligence*). There are some qualifications to be made with regard to subtypes of state adjectives, as will be seen below. However, I will proceed to adopt the theoretical assumption that adjectives denoting states can have independent discourse referent status on the d-axis. Then, *get*-constructions like (18) can be represented within the DST framework:

- (18) a. John got rich, drunk [etc.]
 b. John got tired, angry [etc.]

These examples indicate a distinction between property words (adjectives and past participles) that admit an intentional agentive reading and those that do not. Examples in (18a) allow two interpretations: either John had an intention to become rich, drunk or informed, or the change of state occurred without his agency. In (18b) only the second kind of interpretation seems to be readily available.

Furthermore, not all property words can occur in the *get* construction. The distinction between “stage-level” and “individual-level” states (Carlson 1980) is relevant here.

- (19) ?John got intelligent, generous [etc.]

An explanation for this pattern is that property words which cannot occur with *get* denote properties that are taken to be permanent or inherent to the individual. Some properties

18. English *become*, denoting change of state, developed from Germanic **bikweman*, which also is the source of Modern German *bekommen* “get”. This shows that a motion verb developed both an obtaining sense and a change-of-state sense – diachronic evidence for the conceptual relatedness of motion, obtaining and change of state.




Figure 9. John got rich

are “mobile” and some are not. An agentive reading of (19) requires the reader to invent a facetious context. A word such as *tall* is acceptable with *get*, depending on context. In *John got tall* it may be acceptable if John is a child but not if he is an adult.

DST can provide a model for the agentive and the non-agentive readings by means of a straightforward extension of the basic prehension schema, allowing property words as discourse referents. Thus, Figures 2 and 3 give the required configuration with the following differences, which are incorporated in Figure 9. For agentive readings of (18a), Figure 3 would have, for example, *rich* instead of *apple*. The vector set at t_j in Figure 9 represents the state “John be rich”, etc., or John’s “having the property” of being rich. In this sense the $\{h\}$ set remains conceptually appropriate, since it represents location. In contrast to Figure 3, we have to make use of a feature of the model introduced earlier and crucial for the *get* constructions discussed in the following sections. In order to represent (18a), the vector set at t_i appears not in the realis but in the irrealis plane defined on the m -axis – as in Figure 9. For the non-agentive reading of (18a) and for (18b), r_1 is absent, as in Figures 2 and 6. In this case, John is, as it were, the recipient or the arrival point for the property.

The diagram has to be understood as r_1 and/or r_2 moving the relation *John-rich* from irrealis plane to realis plane. Here “rich” is treated as having a location in the discourse space at $m=0$, i.e. all properties are treated as real. However, their predication of a particular individual may *not* be real. In (18a), the *get* construction includes the meaning that John was not rich at some time prior to the conclusion of the *get* process. At the end of the process, i.e. at $t_n > t_k$, John possesses the property – i.e. it is located at John, analogously to the $\{h\}$ relation in Figure 2. The vector r_1 has to be interpreted as John’s agentive

application of energy on the counterfactual state which results in its progressive realisation over time at t_k . At some time preceding t_j , it is not the case that John is rich; from t_j to t_k it *becomes* the case that he is rich – a kind of emergence over time that the model can capture graphically. The two possible readings of (18a) are automatically accounted for by the initial assumptions about the shape of the base image schema. For the agentive reading, r_1 remains in the picture; for the non-agentive reading (e.g. “she got lucky”) r_2 alone is present. The resultant state encoded by “John is rich” is captured by the set $\{h\}$, interpreted as location, as is also the case for *have*: the property is located at John.

There are two other constructions closely related to that illustrated in (18). The first is *get* NP-Adj/V_{past part}. It follows, once properties are given coordinates, that we can deal with resultatives like (20):

- (20) Mary got John drunk/the window open/the floor clean [etc.]

Second, examples like (18) are closely related to “passive-*get*”, the focus of much comment in the literature, which is exemplified in (21):

- (21) a. John got informed about the situation
b. John got blamed for the situation

While the examples in (18) focus on transient properties construable as states, with no indication of the event causing the state, causative *get* examples such as (21) indicate states (the state of being informed or blamed) but in addition indicate the event that caused the state (cf. Fleisher 2006: 234–236).

In parallel with (18a) and (18b), (21a) attracts agentive reading, while (21b) does not. That is, in (21a) John may or may not have actively sought information, while (21b) typically does not imply actively seeking blame. Reflexive forms of both sentences make the agentive reading obligatory: *John got himself informed*, *John got himself blamed*.

Diagrammatically, these facts can be modelled as outlined in Figure 10. This configuration is derived very simply from Figure 9. While the stative property *rich* was treated as a discourse referent, verbal predicates (in this example the verb *inform* or *blame*) is treated in DST as a directed relation represented as a unit vector, with agents at the tail and patients at the tip, roughly speaking. In (21) we have a transitive verb whose patient is John. There is no expressed agent, but an agent is certainly in the frame, so the model assigns an unlabelled coordinate on the d-axis. Assigning a label – say “the Central Office” in place of the question mark in Figure 10 – would give a model corresponding to the full passive form *John got informed by the Central Office*.

In addition, it is possible to get a reading in which *informed* and *blamed* are stative properties (cf. the remarks of Fleisher (2006: 232–234) on the past participle *acquainted*). In this case, *informed* and *blamed* would be located on the d-axis, as is the case for *rich* in Figure 9.

The agentive and non-agentive readings are captured in the usual way: for the agentive reading both r_1 and r_2 appear; for non-agentive or “receptive” reading r_2 alone remains. As for the oft noted fact that, in contrast to *be* passives, *get* passives frequently seem to lay responsibility on the entity denoted by the grammatical subject – e.g. *John got injured* versus *John was injured* – this can be explained in cognitive terms as a residual effect of




Figure 10. John got informed

derivation from the $r_1 + r_2$ schema, where r_1 implies intentionality, as was argued in more detail above (Section 5). As also noted earlier, such readings are plainer when a reflexive is present: *John got himself injured*. The reflexive *himself* is already implicit in Figure 10: it is any point with the *John* coordinate on the *d*-axis and it can be made explicit by labelling the point where the head of r_1 contacts the head of the event vector (*informed*, *blamed*). This account is plainly different from e.g. Haegeman (1985) and Givón and Yang (1994) in being primarily cognitive. Synchronically, it can be said that *John got (himself) blamed* provides the option of making explicit a conceptual element, denoted by the reflexive pronoun, which is already present; its presence explains, moreover, the residual sense of “responsibility” for the event undergone.

8. Caused events with *get*: *get-NP-to-V-NP*, *get-NP-to-V*

Sentences of the (22) type can all be captured by minor adjustments in the model outlined in Figures 9 and 10.

- (22) a. John got the car mended (by Fred)
 b. The car got mended (by Fred)
 c. Fred got John drunk
 d. John got Fred to mend the car

In these cases the patient of the embedded verb is not the matrix subject as it is in (21) but another discourse entity – e.g. the car in (22a) and (22b), and John in (22c), while the car is the patient in the *to*-clause in (22d). In some sense the constructions exemplified in sentences (22a), (22b) and (22d) make such entities the patient of two forces: of the r_1 component of *get* (for the appropriate reading) and also of the verb *mend*. (22c) brings about a co-location of John and the stative property *drunk*, in similar fashion to (18a, 18b) and Figure 9. Sentence (22b) calls forth a reception-only meaning, while (22a) is ambiguous. Once again, this is accounted for by the cognitive backgrounding (or weaker activation) of the r_1 phase of the base image schema, which is captured as in other cases in the framework.

The DST model assigns a realis coordinate to *the car* or *John*, as in Figure 11a, which models (22d) and Figure 11b, which models (22a). The appearance of *to* in (22d) is explicable (cf. Gronemeyer 1999, 2001) in quasi-spatial terms, analogously to the truly spatial *John got the book to Mary* in Figure 5. This motivation turns out also to be important in explanations of other *get* constructions (see below, Section 10). Additionally, from Figure 11b we can derive (22b), as we can derive *Mary got the box* from *John got the box to Mary*. Furthermore, we can derive *Fred got to mend the car* from Figure 11a. However, the interpretation of this latter type of sentence is debatable. It has been claimed (Gronemeyer 1999, 2001) that such a sentence can take a deontic (either permission or obligation) reading. Such a reading is not determined by the model, which predicts a reception reading only. This issue is discussed in further detail below (again, see Section 10 below).

It is worth noticing that the base schema (see Figures 1b and 3), which includes an outcome “possession” or “having” phase, also transfers to models of (22a, b, d) and perhaps (22c). As often observed, causative *have* is very close to causative *get*, and given the assumption of the base image schema, it is possible to see why this should be so. However, it is also sometimes noted that *get*-causation seems to imply an indirectness that *have* causation does not. This is clearer in examples analogous to *John got Fred sacked* versus *John had Fred sacked*, where the latter sentence seems to imply that John was the more direct agent of the sacking. If we include the {h} vector in the diagram, as in Figures 11a and 11b, we do indeed get a direct relation between John and the respective patient entities, as entailed by the base model.

The modal dimension provided in the present theoretical framework is important for the analysis of *get*. While Fred and the car are realis, the event that we can represent propositionally as *mend* (f, c) is not. It is either counterfactual or merely possible from S's viewpoint. It is shown here in Figures 11a and b as counterfactual. Vector r_1 puts pressure, so to speak, on Fred – or his counterpart in the irrealis plane, i.e. on his coordinate. Vector r_2 is the realisation phase in which the event is completed over time, that is, “real-ised”, and is the derivative of the reception vector in the basic schema. Thus reception, the obtaining of an object, is extended conceptually to the realisation of an irrealis event intended, phenomenologically speaking, by a subject. It should be noted that the *have*-causative construction does not involve a *to*-infinitive. The *have*-construction backgrounds the prior “getting” event (though it is available) and foregrounds the result, i.e. “having”.¹⁹

19. The account presented in this section is an alternative to one that would simply show r_1 causing the event of mending in the realis plane, perhaps treating Fred as being “moved” from one activity to another,




Figure 11a. John got Fred to mend the car




Figure 11b. John got the car mended by Fred

Several further semantic phenomena associated with the constructions surveyed in this section need to be considered. Consider, for example, the following contextualisations for (22a):

- (23) a. John finally got the car mended – he'd been putting the task off for weeks
 b. John got his car mended by Fred, after offering him an extra ten quid
 c. John got his car mended by the great *Fred*, lucky man

All three readings can probably be accommodated in the present framework.

There are also causatives of the following form *get*-NP-Ving/Ved, *get*-Ving, *get*-Ved:

- (24) John got Fred talking, the machine working/*worked
 John got Fred started/*starting, the machine started/*starting
 John got Fred starting to play well
 John got working/*worked, started/*starting

The distributional patterns reflect the aspectual semantics of *get* but the details and the appropriate modelling cannot be considered further here.

9. Is *get* to deontically modal?

Gronemeyer (1999 and especially 2001) offers the most detailed account to date of modal uses of *get*. Her analysis of the basic non-modal uses as “ingressive+*be*+Prep” seems to be broadly compatible with the DST account outlined in the present paper. Gronemeyer’s account of modal *get* is compatible to a degree, but there are some serious problems with her 1991 and 2001 exposition.

One difficulty with Gronemeyer (1999, 2001) concerns interpretation of the data. She proposes that *get* has two modal meanings, permission and obligation (repeated by Manna 2004: 3). The following example is used by Gronemeyer (2001: 5):

as suggested by Hollmann (2003). Gronemeyer (1999:25), for example, analyses the *get*-NP-*to*-V_{inf}-NP construction as analogous to *get*-NP-[*to*-NP]_{pp}. That is, the complement of *to* is type-shifted, from NP to VP. Hollmann (2003: 99–102) finds historical evidence against Gronemeyer’s proposal and offers an alternative account. Hollmann also, however, rejects some analyses using the obtaining notion on the grounds (p. 106) that it is stretching the “obtaining” concept too far. Hollmann’s account uses the idea of a schema abstracted from usage occurrences, to account for the *get*-NP1-*to*-NP2 construction developing into *get*-NP-*to*-VP. This schema he notates as *get* NP-*get*-NP-*to*-XP. My account is somewhat different from all of these approaches in that it involves an explicit modal dimension, although the importance of this dimension is hinted at in Gronemeyer. However, my account does retain a quasi-spatial movement concept that may indeed be linked with the presence of *to*, namely the “movement”, via the *get* operator, of an irrealis event into the realis plane. It is not the actor that is “moved” but the conceptualized event. The problem with viewing the *get* construction in question as “moving” an actor to a new location or to a new activity is that activities and the acquisition of properties cannot be imagined as existing referents prior to their occurrence in the way that locations can – this is one reason why the modal axis has to be invoked. Essentially, the modal dimension deals with the widely acknowledged inchoativeness of *get* constructions.

- (25) a. You get to have all the fun,
 b. while I get to do the dirty work

Gronemeyer thinks that the default for uses like (25a) is permission but that this reading can be overridden if the embedded VP expresses an action with negative evaluation. Consequently, she claims that (25b) expresses obligation. While she is surely correct to point out that the meaning of the complement (and other contextual factors) is crucial, her reading of these constructions is not entirely convincing. She claims that sentences like (25a) express permission, at least in American English, on the basis of native-speaker intuition and anecdotal evidence (1999:7). I think this assertion needs further analysis, certainly for other varieties. In my variety of British English it is not the case that (25a) unequivocally expresses permission as such. The claim that (25b) expresses obligation is open to questioning, on my own native-speaker intuitions. The following evidence is relevant to the interpretation of both sentences.

It may be true that one can grant John permission by uttering: “OK, John gets to go to the party” in the sense of “John may go to the party”. But so may other expressions that do not use modal verbs. More contemporary corpus data are needed to establish whether this usage is becoming entrenched for performing the speech act in question. Even if there is evidence that *get to* is moving towards grammaticalisation as a deontic modal, in current English there remains evidence of a semantic distinction:

- (26) a. ?Jane got to go to the party but she decided not to
 b. Jane was allowed/permitted to go to the party but decided not to

(26a) seems unacceptable because *got to* entails the subject did go: *get* is an achievement operator. On the other hand (26b) expresses the possibility for Jane to go; however, possibilities need not be realised. Haegeman (1985) rehearses the syntactic reasons for thinking that *get* is still a lexical verb, not a grammaticalised modal, in English contemporaneous with her time of writing.

Further, (25b) is in my view mis-analysed by Gronemeyer, who claims *get* is coerced into an obligation meaning by the negatively valued complement *do the dirty work*. In my reading, *get* in (25b) retains its semantics and what the semantics of the complement does is induce a pragmatic effect: the reader interprets the sentence as an ironic utterance under the influence of the expression *dirty work*.

These observations suggest that the paraphrase produced by Hoekstra (1994) and followed by Gronemeyer (2001) is not appropriate: “if one has permission to do something, one has received the possibility of doing it” (Gronemeyer 2001:5; cf. 1999:7, 31). But (26a) and (26b) do not seem to entirely bear this out. If *get* in (26a) expresses permission, it should on this definition express (mere) possibility, whereas in fact there is an entailment that the event of Jane’s going to the party did in fact occur. This is distinct from (26b) where we see genuine possibility.

Gronemeyer derives modal *get*-constructions from causatives, essentially by considering pairs such as the following:

- (27) a. John gets Mary a book ... Mary gets a book
 b. John gets Mary to leave ... Mary gets to leave

In syntactic terminology the second sentence of each pair has an external argument (subject) derived transformationally from an internal argument in the first pair. Informally, “the subject makes the recipient have/do XP” (Gronemeyer 1999: 31). In (27) *Mary is a recipient*. This may be true, and is consistent with the model developed in this paper. However, it does not at all explain the semantic facts. Gronemeyer (1999: 31) claims that *get to is*

now generally modal and is currently able to denote obligation as well [as permission]; which end of the deontic scale is expressed is determined contextually.

This is implausible, as we have seen, given the doubts about the allegedly obligation meanings of sentences like (25b). There is little evidence that *get to* has (yet) developed an intrinsic deontic meaning, at either end of the scale. To explore the possibility of contextual determination we can consider sentences in which the lexical material in the complement of *get* invokes conventionally deontic frames. For example, consider:

- (28) a. The prisoner gets to exercise in the yard
 b. The prisoner gets to stay in his cell twenty-four hours a day

Though the context of (28b) demands an obligation meaning for the matrix verb, *get* does not appear to me to supply it. I have to read (28b) as an ironic utterance. In (28a) the interpretation could indeed be permission, but is cancelled if one adds “when the gaoler is not looking”, something that would not be the case if *get to* actually *entailed* permission. This is not to say, however, that somehow a permission reading is not available, and conceivably facilitated by the meaning of *get to*. However, to treat *get to* as having an “established” obligation reading seems mistaken.

How should we, then, characterise the meanings of *get to*? In terms of the present framework, *get to* meaning constructions are motivated by the prototype construction, by way of the following derived constructions analysed above:

- (29) a. John got Fred to drive the Bentley
 b. John got to drive the Bentley

As was mentioned in section 8 above, a sentence like *Fred got to mend the car* can be derived from the causative *get* construction. One might expect the causative element, represented by the force vector, to yield an obligation meaning, but as argued above the meanings of this sentence type do not seem to warrant such a conclusion. An alternative motivation lies in sentences of the motion *get* type: *John got to London* What we have also seen is that *to* can be explained in terms of “modal movement” – that is from the irrealis plane into the realis plane in DST terms (cf. also Duffley 1992; ter Meulen 1995 and Gronemeyer 1999, 2001).

In (29a) there is force on Fred from John represented by r_1 . In (29b), if there is any possible reading involving willed effort on the part of John, it is residually derived from the base image schema; however, the default reading is that John “received” and achieved the embedded event. The concept of applying force to oneself to receive something is self-contradictory: hence it is also odd to say “John got himself to drive the Bentley”. The sense of the *get to* construction is further motivated by the sense (already analysed) in the

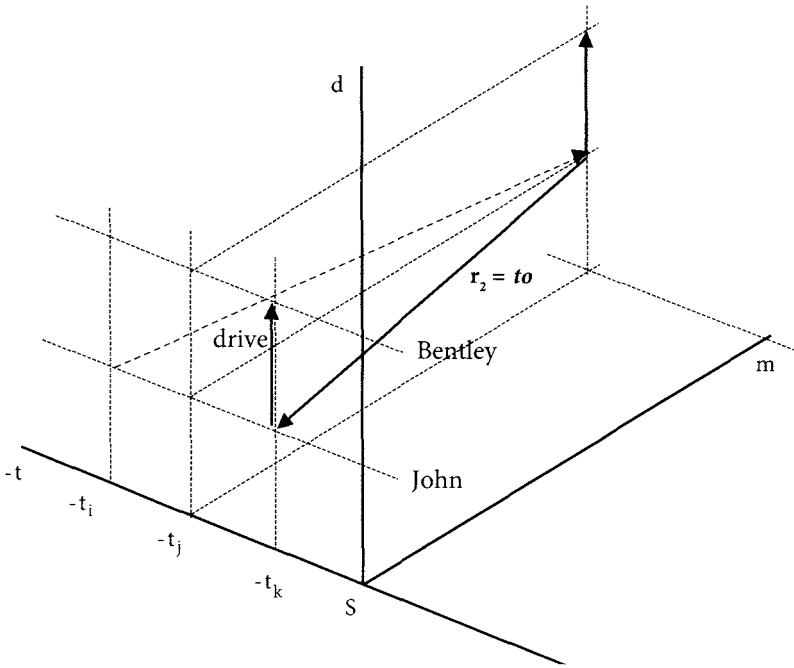


Figure 12. John got to drive the Bentley

following: the prototypical cases like *John got the apple* in the reception-only (the r_2 only) sense (see Figure 3), *John got a medal* (see example (6b)), and *John got rich* (see example (18a) and Figure 9). To get a permission reading out of (29b) it is of course possible to invent an appropriate pragmatic context, but this is not the default interpretation.

Figure 12 should be understood as “moving” an irrealis event to the realis plane – as “real-ising” it. The event does not depend on the effort of John, or the application of some sort of force by him. It is possible, however, given appropriate contexts to have an interpretation of (29b) in which John does in fact make wilful effort to bring about the result: Figure 12 indicates this possible reading by the dashed r_1 force vector. While a permission context is also easy to supply, it is not an intrinsic part of the meaning of this construction and so finds no expression in the figure.

10. Obligation: *have got NP to-V*, *have got to-VP*

Most accounts of the modality of *have got to* do not get much further than describing the syntax (limitations on tense to present and present perfect), noting the underlying possession senses and relating *have got to* to *have to* (Gronemeyer 1999, 2001; Manna 2004). Gronemeyer’s explanation is almost entirely diachronic and purely syntactic, deriving modal *have got to-VP* historically from *have got NP to-VP*, driven by word order change in Middle English (based on Fischer 1994). For Gronemeyer, the meaning shift involved in *have got to* is explained as being based on “the inference that if you have something to

do, then you are obliged to do it" (Gronemeyer 2001: 33, citing Bybee et al. 1994: 181–187). Unfortunately, this debatable formulation does not seem per se to explain the conceptual shift between possession and obligation: indeed, as it stands, it appears to be almost circular. There is nothing in the existing accounts that explicitly models the concept of obligation in such a way as to show how it is related to that of having or possession. I propose a very tentative DST analysis of the deontic *have got to* construction that seeks to relate this construction to the base schema via the derived schemas outlined in preceding sections.

First, let us review a few of the peculiarities of the *have got to* construction. In British English (30a) is ambiguous:

- (30) a. John has got to drive the Bentley
b. John has gotten to drive the Bentley

The past participle form in (30b), mainly typical of American English, does not include the modal meaning. This points to the special semantic status of *have got to*.

It is widely noted that *have got*, both modal and possessive, is related to *have* (see also Section 4 above) and that the modal *have got to* construction is close to both *have to* and *is to*.

- (31) a. John has to pay the fine
b. John has got to pay the fine
c. John is to pay the fine

This is potentially relevant, since the common ground appears to involve a locational concept. Also *have* and *have got* are evidently not fully synonymous and one way to explain this is to claim that *have got* retains conceptual links with the full schema of *get*, outlined in Figure 1b.

Both the *have* and the *have got to* construction appear in the following variations:

- (32) a. John has a report to write [agent force, or weaker force]
b. John has to write a report [external force/agent force, or stronger force]
c. John has got a report to write [agent force, or weaker force]
d. John has got to write a report [external force/agent force, even stronger force]

As already noted, *have got* whether modal or possessive, is tense-restricted, to either present tense or present perfect. At first sight modal *have got* appears to be equivalent to possessive *have got*. This is not quite the case. In its possessive sense, the past tense form of *have got* is questionable:

- (33) a. He's got blue eyes/he has blue eyes
b. ?He'd got blue eyes/he had blue eyes

In its modal sense *have got* accepts both tenses, at least in some varieties:

- (34) a. He's got to pay the fine
b. He'd got to pay the fine.

Further clues may arise from negation. Modal *have got to* is distinct from modal auxiliaries in that negation scopes over the modal meaning and not over the meaning of the embedded *to*-infinitive:

- (35) You must not drive on the right in Britain; you must drive on the left
 You haven't got to drive on the right in Britain; you've got to drive on the left
 ?You have got not to drive on the right in Britain
 *You don't must observe the rules here (you haven't got to observe the rules here)
 *You don't got/get to drive on the right
 *You don't have got to drive on the right

Negation has the effect of denying the obligation. The negative particle attaches to *have*, which thus appears to be acting like an auxiliary to a head *got*, albeit not of the usual *do* form. Some varieties appear to do without the verbal auxiliary, although the evidence is not unequivocal:

- (36) We got to do it
 We gotta do it
 *We gotta not do it
 *We don't gotta do it,

where assimilation is an indication of grammaticalisation. These considerations seem to indicate that *have got to* may be in a state of change diachronically. Importantly, it seems to have not lost its conceptual link with basic meanings of *get*, while having acquired some but not all of the properties of modal verbs.

Two semantic peculiarities need also to be mentioned. First, *have got to* is, like *must*, modally strong:

- (37) ?John has got to pay the fine but he won't
 ?John has got to pay the fine but he might not
 ?John must pay the fine but he won't
 ?John must pay the fine but he might not
 John should/ought to pay the fine but he won't
 John should/ought to pay the fine but he might not

In all cases it is presupposed that John has not yet, at utterance time, paid the fine. Secondly, both high-probability assertion about the future and mid-probability assertion may be incompatible conceptually with *must* and *have got to*. Thirdly, it seems to be the case that, in sentences with first-person grammatical subject, *have got to* deontic meanings are not typically taken to be speech acts performed by the speaker; rather, they seem to denote other-imposed obligations. Deontic *must*, on the other hand, seems to be compatible with both situations.

The crucial features to emerge from the above overview seem to be: the importance of the present tense restriction; the idiosyncratic combination of realis and irrealis elements of meaning; the appearance of a source other than the speaker; and the pervasive presence of the base schema. I proceed now to outline a model of the *have got to* construction consistent with these points and with the framework developed so far. It should, however, be emphasised that this proposal is provisional.

On the assumption that both *X has got Y* and *X has got to Y* have cognitive links to the full schema (Figure 1b), Figure 13 offers a partial DST model of these constructions. The most obvious point to note is that the geometrical configuration of Figure 13 is essentially the same as that for the causative *get*-constructions discussed in Sections 7 and 8 (Figures 9 to 11). This corresponds to the intuition that there is some causative element in the concept of obligation, in the sense that this concept implies that some socially based authority exerts some kind of causal force, though achievement is not asserted or entailed. The purpose of the exerted force is to cause some agent to bring some event or state into reality.

Syntactically, the complement of *has got* is a clause headed by *to*. This clause is an event that is, as in other cases examined, transferred from the counterfactual plane to the realis plane. In example (34a), John has the event imposed upon him: he “receives” this event. We cannot of course say that the paying of the fine is located at S’s utterance time in the present-tense version (34a). Although “has got” does include S’s utterance time, “to pay a fine” does not. This is because present “now” is a point, as the geometry of the discourse space model makes clear, and events are not punctual (Michaelis 1998; Chilton 2007). Consequently, in deontic contexts, if the complement of *has got* is an event, the time-reference is coerced into the future, where there can be spans of time. In epistemic contexts the time reference of event complements can include the present, but their form has to be present continuous, as for any eventive as opposed to stative verb. However, if the complement of *has got* is a stative verb, then there is future time reference only if the contextual interpretation is deontic, while if the reading is epistemic it is generally taken to coincide with the present. Consider in this regard *John has got to be law-abiding*, which can be ambiguous between deontic and epistemic readings. The point here, in relation to our treatment of deontic *has got to+V* is that this construction operates on a counterfactual eventuality (state or event) denoted by the verb and moves it into the realis plane, whence it may be coerced into a future time location. All future eventualities are inherently epistemically uncertain.

These points are consistent with a conceptual analysis of the particle *to* with English infinitives. *To* occurs generally with verbs in clauses that are complements of psychological verbs such as *hope*, *want*, *forget*, *remember*, etc. (cf. Langacker 1991:438–463). The basic spatial meaning of *to*, roughly “movement of A in direction of goal b”, has acquired an abstract grammaticalised meaning on the basis of the PATH schema, by a metaphorical process yielding something like “future intended actions are spatial goals”. In general, actions or events, denoted by clauses under the scope of *to*, are future relative to the individual denoted by the grammatical subject of the matrix verb. This pattern combines with *has got* in the construction we are considering.

Figure 13 shows only the main effect of the modal *have got* operator. While the core configuration remains structurally similar to the base *get* schema, and to the models for the causative *get* constructions, there are specific differences related to the specific meaning of the construction. The most important feature is the foregrounding of r_2 : the grammatical subject of modal *have got* is not the agent (not the source of r_1) but the recipient (the goal of r_2). In (31b) John is the “recipient” not the getter. The forerunner of this kind of meaning are cases such as “John has got a medal” (cf. the discussion of (3) and (4) in




Figure 13. John has got to pay the fine

Section 3). What of the force vector r_1 ? It is possible to model the modal *have got* instances by leaving out r_1 and there may be pragmatic contexts in which this is the appropriate model. However, as argued for other *get*-constructions, r_1 is never entirely absent and for many contexts of the *have got to* construction some kind of r_1 is probably activated. Since vectors have to have some source coordinate, in Figure 13 there is a source of the force vector r_1 on the d-axis, but this coordinate is not specified (labelled). This force vector, inherited from the base *get* schema, is important for the emergence of the deontic meaning, and can be interpreted as *deontic force*, in line with the Talmy-Sweetser force-dynamic explanation of modals (Talmy 1988; Sweetser 1990).

As in the causative *get*-constructions discussed in Sections 7 and 8, r_1 results in translation of an irrealis event (*John pay the fine* in (31b)) to the realis plane. Figure 13 shows the result of coercion of the event (*pay'*) into what is, relative to S's discourse space, an epistemically uncertain future, located diagrammatically at the m-axis mid-point and at an arbitrary point on the t-axis $t_1 > t_0$. The speaker of (34b) cannot be sure that John will pay the fine, only that he "has got" (the obligation) "to pay the fine" in the future. In this particular diagram, "the fine" itself is treated as realis, i.e. its m-coordinate is 0. In the discourse space it has a distinct existence as an element in a real legal frame, but at utterance time the paying of it is not real. The sense in which the paying is not real is not the same as the sense in which the source of r_1 is not "specified" (or labelled). There is a real source of r_1 but its source in any particular case has to be inferred by the hearer, by the

use of contextual information: for example, the source of moral force might be “the rules of the game”, “custom”, “the law”, “God”, etc. On an alternative reading in which it is the speaker who is the source of the deontic force, the tail of r_1 has a coordinate at $d = 0$, i.e. at S. This is the case where, for example, the speaker of (31b) is a magistrate in a court.

11. Conclusion

This paper has attempted to elucidate the complex semantics of *get* within an analytical framework that is entirely different from other accounts but which is, I would claim, consistent with cognitive-linguistic approaches. The premise is that word meanings derive from frames and in many cases from image schemas. In the case of polysemy, variant but related meanings are also related to a frame or to an image schema. To account for the polysemous *get*, I postulated that the word's basic meaning is associated with an image schema that I have called prehension, a supposition justified by proto-forms and attested early meanings. Further, however, I have utilised the DST method of formally representing conceptual structure, based on the proposal that many core meanings engage an abstract “discourse space” defined on three dimensions. This approach involves the further assumption that directionality is conceptually central; this assumption turns out to be important for the formal modelling of the semantics of *get*. Using these ingredients I formalised the prehension schema – parameterized it – by means of vector components. In brief the idea is that basic *get* is a schema that consists of (i) a force (or effort) vector (r_1) directed away from the self or agent towards an object and (ii) a vector (r_2) that represents the movement of the object in the direction of the self or agent. The object is then located (received, had, possessed) by the self or agent.

From this base schema we have explored how the other meanings of *get* may be derived. These extensions involve simple operations. It seems that the first component vector, the force vector, may drop out or may be only weakly activated. The notion of “weak activation” was introduced to account for the often reported sense that some *get* constructions, even when they do not seem to express intentional agency, nevertheless seem to imply responsibility or fault. The second vector, however, the reception component, does not seem to drop out in any meaning. Directionality may change. End-state location (possession represented as {h}) may not be foregrounded, i.e. may be only weakly activated.

The first natural extension discussed (Section 5) simply moves the receiving location from the self/agent to some distal recipient or location. The result is a linear path – in effect a path schema – consisting of the two component vectors, the force (or effort) vector and the receiving vector. Conceptually, in Aktionsart terminology, this gives the schema underlying for achievement verbs. In this manifestation of *get*, the effort component may or may not be present, i.e. it may be weakly activated, depending on context. This account gives a motivated way of understanding how *get* comes to mean movement of the referent of the grammatical subject. Basically, *get* schematizes effortful action on an object that results in the relocation of the object at self; then the same schema is extended to effortful action resulting in movement of object to other or distal point, which results in a path schema. This path schema then appears to be available as a representation of the move-

ment of an agent (and also inanimate entity) to a distal point, with intentional agency maybe activated by context, maybe not (Section 6). In Section 7 I proposed that properties could be treated as entities that change location and come to be possessed – interestingly, the *property* – of the grammatical subject. Again, the effort vector component of the base schema may or may not be fully activated.

There is a crucial further element for this and the other constructions discussed from Section 7 onwards – the dimension of modality. In the case of *get*+property constructions, there is a condition in which effort is transmitted over time onto an abstract relation between self and object (property) in the counterfactual plane – i.e. there is a time in the represented event at which the grammatical subject referent does *not* possess some property and a later time when it does. By contrast, the object in the basic *get* schema exists the whole time. This is an important point because it means that our descriptive apparatus must have the means to formalise degrees of reality, as DST does. Moreover, the same requirement seems to be present for the causative and other constructions discussed following Section 8. In the case of causative *get* constructions, events also cannot be said to exist until they happen: they are “brought” into being over time – something that the formalism used here can easily capture.

Historically speaking, the latest development of *get* meanings seems to be its deontic meaning. It is possible that we can observe change in process in the differences between English varieties and the doubts raised in Section 9. In certain varieties there is a construction that has a “success” or “achievement” reading that is inherent in the basic schema and arises when the effort vector is absent or weak, given certain contexts. What is not in doubt is that the fixed form *have got to V* has a strong deontic meaning. At first glance, it would seem to be a conceptual jump from prehension to normative obligation – a puzzle that surely deserves some speculative explanation. In the account sketched in Section 10, the deontic meaning falls out from the basic schema, though there are many aspects of this approach that need further investigation and refinement. The fundamental configuration remains: two components and a resulting “relocation” of an “object.” As in the cases noted earlier, an irrealis event is relocated in the conceptual space – specifically, for this construction, into the modally positive present, the point t_0 . The future is subject to the conceptual characterisation that it may or may not occur, but given this restriction, *has got to* transfers an irrealis event into the strongest possible reality in future conceptual space. There is a further important but simple modification of the geometry of the base schema. Just as we saw the possibility of a shift of the receiving point (the end point of r_2) from self to distal point, so in the *has got to V* construction, it is the start point of the force vector that is shifted to a distal point, interpreted as the source of force. In this case we can think of r_1 as representing “moral force,” its source location and its impact location, and thus as consistent with the Talmy-Sweetser use of force dynamics to model modality.

The most detailed recent studies of *get* have focused primarily on diachronic syntactic derivations of *get* constructions in order to explain synchronic polysemy. Fleisher (2006) emphasises that semantic motivation has also to be coordinated with such accounts and in practice Gronemeyer also includes some semantic motivation. However, since the generativist background separates semantics and syntax, this point has to be re-made by the analysts using it. In the CL framework grammatical form and conceptual meaning are

intrinsically linked: any account of grammatical form should also be an account of conceptual structure. In the present study of *get* constructions I have not attempted to give a diachronic account. However, I propose that the account given here of conceptual-grammatical relations between the *get* constructions provides the conceptual motivation for diachronic change.

An implicit claim in the present account is that *get* retains its basic image-schema structure in current English constructions. There are no (as yet) fully grammaticalised meanings. The ambiguity reading of the sentence with which we started out – relating to whether the subject is interpreted as acting in a goal-directed fashion – might be taken as evidence that the full schema is indeed active. The central claim, however, is that the extraordinary range of associated meanings in English *get* can be explained as produced from the prehension image schema. This verb is a classic example of embodiment, in which motor-spatial representation provides the ground for polysemy.

References

- Bonnefille, S. 2006. Constructions with *get*: How to get the picture without getting confused. *Annual Review of Cognitive Linguistics* 4: 21–37.
- Bybee, J., Perkins, R. & Pagliuca, W. 1994. *The Evolution of Grammar: Tense, Aspect and Modality in the Languages of the World*. Chicago: University of Chicago Press.
- Carlson, G. N. 1980. *Reference to Kinds in English*. New York: Garland Publishing.
- Chilton, P. 2005. Discourse Space Theory: Geometry, brain and shifting viewpoints. *Annual Review of Cognitive Linguistics* 3: 78–116.
- Chilton, P. 2007. Geometrical concepts at the interface of formal and cognitive models: *Aktionsart*, aspect and the English progressive. *Pragmatics and Cognition* 15 (1): 91–114.
- Clausner T. C. & Croft, W. 1999. Domains and image schemas. *Cognitive Linguistics* 10 (1): 1–31.
- Croft, W. 2000. *Verbs: Aspect and Argument Structure*. Work in preparation under contract to Oxford University Press, <http://www.unm.edu/~wcroft/Verbs6.pdf> accessed 25 May 2007.
- Croft, W. & Cruse, A. 2004. *Cognitive Linguistics*. Cambridge: Cambridge University Press.
- Duffley, P. 1992. *The English Infinitive*. London: Longman.
- Feldman, J. A. 2006. *From Molecule to Metaphor: A Neural Theory of Language*. Cambridge MA: MIT Press.
- Fillmore, C. 1982. Frame semantics. In *Linguistics in the Morning Calm*, The Linguistic Society of Korea (ed.), 111–137. Seoul: Hanshin.
- Fillmore, C. 1985. Frames and the semantics of understanding. *Quaderni di Semantica* 6: 222–254.
- Fischer, O. 1994. The development of quasi-auxiliaries in English and changes in word order. *Neophilologus* 78: 137–64.
- Fleisher, N. 2006. The origin of passive *get*. *English Language and Linguistics* 10 (2): 225–252.
- Gallese, V. 2003. A neuroscientific account of concepts: From control to representation. *Philosophical Transactions of the Royal Society of London* 358: 1231–1240, published online 4 June 2003.
- Gallese, V. & Lakoff, G. 2005. The brain's concepts: The role of the sensory-motor system in conceptual knowledge. *Cognitive Neurophysiology* 21 (3–4): 455–479.
- Givón, T. & Yang, L. 1994. The rise of the English *get*-passive. In *Voice*, B. Fox & P. J. Hopper (eds), 119–149. Amsterdam: John Benjamins.
- Goldberg, A. 1995. *Constructions*. Chicago: University of Chicago Press.
- Goodale, M. A. & Milner, A. D. 1996. *The Visual Brain in Action*. Oxford: Oxford University Press.

- Gronemeyer, C. 1999. On deriving complex polysemy: The grammaticalization of *get*. *English Language and Linguistics* 3(1): 1–39.
- Gronemeyer, C. 2001. Modal readings of light verbs with to-infinitivals. *Working Papers* 48. Lund University, Department of Linguistics, <http://ask.lub.lu.se/archive/00019162/Gronemeyer.pdf>, accessed 27 March 2007.
- Haegemann, L. 1985. The *get*-passive and Burzio's generalization. *Lingua* 66: 53–77.
- Hoekstra, T. 1994. HAVE as BE plus or minus. In *Paths towards Universal Grammar: Studies in Honor of Richard S Kayne*, G. Cinque, J. Koster, J.-Y. Pollock, L. Rizzi & R. Zanuttini (eds), 199–215. Washington: Georgetown University Press.
- Hollmann, W. 2003. Synchrony and Diachrony of English Periphrastic Causatives: A Cognitive Perspective. PhD Dissertation, University of Manchester.
- Holmes, N. P. & Spence, C. 2004. The body schema and the multisensory representation(s) of peripersonal space. *Cognitive Processes* 5 (2): 94–105.
- Johnson, M. 1987. *The Body in the Mind*. Chicago: University of Chicago Press.
- Lakoff, G. & Johnson, M. 1980. *Metaphors We Live By*. Chicago: University of Chicago Press.
- Lakoff, G. & Johnson, M. 1999. *Philosophy in the Flesh*. New York: Basic Books.
- Langacker, R. 1987. *Foundations of Cognitive Grammar*, volume 1. Stanford: Stanford University Press.
- Langacker, R. 1991. *Foundations of Cognitive Grammar*, volume 2. Stanford: Stanford University Press.
- Lee, K. 2005. Meanings of *get*: Variations on the same theme. Paper presented at the 8th International Cognitive Linguistics Conference, Seoul.
- McIntyre, A. 2005. The semantic and syntactic decomposition of *get*: An interaction between verb meaning and particle placement. *Journal of Semantics* 2: 401–438.
- Mandler, J. M. 2004. *The Foundations of Mind*. Oxford: Oxford University Press.
- Manna, J. 2004. "Get out of here! Is get even undergoing grammaticalization at all?". Seminar in Syntax: Grammaticalization, Spring 2004, http://www.eden.rutgers.edu/~jgmanna/papers/get_rev.pdf, accessed 4 March 2007
- Michaelis, L. 1998. *Aspectual Grammar and Past-tie Reference*. London: Routledge.
- O'Keefe, J. 1996. The spatial prepositions in English, vector grammar, and the cognitive map theory. In *Language and Space*, P. Bloom, M. Peterson, L. Nadel & M. Garrett (eds), 277–316. Cambridge MA: MIT Press.
- O'Keefe, J. 2003. Vector grammar, places, and the functional role of the spatial prepositions in English. In *Representing Direction in Language and Space*, E. van der Zee & J. Slack (eds), 69–85. Oxford: Oxford University Press.
- O'Keefe, J. & Burgess, N. 1996. Geometric determinants of the place fields of hippocampal neurons. *Nature* 381: 425–428.
- O'Keefe, J. & Nadel, L. 1978. *The Hippocampus as a Cognitive Map*. Oxford: Clarendon Press.
- Piaget, J. 1952. *The Origins of Intelligence in Children*. New York: International University Press.
- Piaget, J. & Inhelder, B. 1971. *Mental Imagery in the Child*, translated P. A. Chilton. London: Routledge.
- Pustejovsky, J. 1988. The geometry of events. *Lexicon Project Working Papers* 24: 19–39. Cambridge MA: MIT Press.
- Rizzolatti, G., Fogassi, L. & Gallese, V. 2000. Cortical mechanisms subserving object grasping and action recognition: a new view on the cortical motor functions. In *Cognitive Neuroscience*, M. S. Gazzaniga (ed.), 539–552. Cambridge MA: MIT Press.
- Rizzolatti, G., Fadiga, L., Fogassi, L. & Gallese, V. 1997. The space around us. *Science*: 277: 190–191.
- Slobin, D. I. 1996. Two ways to travel: Verbs of motion in English and Spanish. In *Essays in semantics*, M. Shibatani & S. Thompson (eds), 195–217. Oxford: Oxford University Press.
- Sweetser, E. 1990. *From Etymology to Pragmatics*. Cambridge: Cambridge University Press.
- Talmy, L. 1985. Lexicalisation patterns: Semantic structure in lexical forms. In *Language Typology and Syntactic Description, Volume 3: Grammatical Categories and the Lexicon*, T. Shopen (ed.), 57–149. Cambridge: Cambridge University Press.

- Talmy, L. 1988. Force dynamics in language and cognition. *Cognitive Science* 12: 49–100.
- ter Meulen, A. 1995. *Representing Time in Natural Language: The Dynamic Interpretation of Tense and Aspect*. Cambridge MA: MIT Press.
- Zwarts, J. 1997. Vectors as relative positions: a compositional semantics of modified PPs. *Journal of Semantics* 14: 57–86.
- Zwarts, J. 2003. Vectors across spatial domains: From place to size, orientation, shape, and parts. In *Representing Direction in Language and Space*, E. van der Zee & J. Slack (eds), 39–68. Oxford: Oxford University Press.
- Zwarts, J. & Winter, Y. 2000. Vector space semantics: A model-theoretic account of locative prepositions. *Journal of Logic, Language and Information* 9 (2): 171–213.

Motion scenarios in cognitive processes

Stéphanie Pourcel

1. Introduction

Linguistic relativity suggests that language patterns and semantic representations influence cognition, so that speakers of different languages entertain differing conceptualisations of otherwise similar events and entities (Whorf 1956; Lucy 1992). In modern cognitive science, linguistic relativity is approached as a hypothesis in need of scientific investigation and empirical evidence. Addressing such a vast question is best achieved by selecting a given domain of experience, investigating its available means of expression in different languages, using the cross-linguistic differences identified as foundations for hypothesising how they may then influence conceptual representations of the selected domain (Lucy 1997).

The chosen domain in this study is space, and more specifically, motion (as opposed to locational reference, e.g. Levinson 2003). Motion is an ideal domain of investigation for several reasons: (i) it pervades human experience and is thus an essential domain of conceptualisation and expression, (ii) it is a complex enough domain not to be amenable to neurophysiological determinism (cf. colour tradition, e.g. Berlin and Kay 1969), (iii) it remains experientially grounded in human and terrestrial physics and hence it may be culturally subjective to a limited extent only (cf. kinship research, e.g. Danziger 2001), (iv) it is expressed differently across languages, including closely-related ones (cf. Talmy 1985), and (v) its means of expression reach beyond lexical resources to the sentence and text levels (cf. Slobin 2004).

This paper begins with an outline of these linguistic differences, with a special focus on English and French. It then addresses methodological issues in relativistic experimentation by reviewing a number of relevant studies with similar aims to the present research, namely to provide supportive evidence for linguistic relativity in the domain of motion (e.g. Gennari et al. 2002; Papafragou et al. 2002; Finkbeiner et al. 2002; Bohnemeyer et al. 2004; Zlatev and David 2004, 2005). These studies fail to meet these aims, as none offers data in favour of relativistic conclusions (but see Oh 2003). At the same time, these also fail to agree in their findings. This chapter offers a suggestive discussion to try and explain this lack of concordance in methodological terms. In addition, this chapter shows that these studies have not considered facts of motion conceptualisation, such as constraints and biases, that exist independently of language. In conclusion, it is argued that the domain of motion remains to be consistently analysed and understood prior to relativistic

applications. This chapter presents a few preliminary studies to cater for an understanding of this type (Kopecka and Pourcel 2005, 2006; Pourcel 2004, 2005). These studies examine the dynamics of motion conceptualisation and suggest that a number of fundamental variables determine the relative cognitive salience of motion dimensions (e.g. *MANNER*, *PATH*), regardless of the cogniser's native language. These variables include figure animacy and humanness, *PATH* telicity, *MANNER* force dynamics, and motion causality.

Based on these considerations, the present study proceeds to testing linguistic relativity in the domain of motion at a more controlled level of methodological and conceptual understanding. The methodology employed uses a motion 'scenario' as stimulus, instead of isolated motion events, and compares the linguistic and cognitive performances of 22 English speakers and 25 French speakers in recall and inferencing tasks. The data demonstrate that the same motion scenario is not only expressed differently by the two language communities, but that its details and sub-events are also memorised and inferred differently in cognition. Importantly, these differences reflect semantic foregrounding and backgrounding for expressing motion in the two languages. The data therefore report a consistent correlation between the semantic framing of motion events in different languages and the conceptual representation of these events in speakers' minds. In short, the data are strongly suggestive of linguistic relativity effects.

2. Motion in language

The domain of motion has been extensively researched in cognitive linguistic typological work (e.g. Aske 1989; Slobin 2004; Talmy 1985, 1991, 2000; Zlatev and Yangklang 2004). Cross-linguistic investigations have revealed that motion events are encoded in language via the mapping of a few central conceptual components, including a moving entity, or *FIGURE*, a spatial reference of displacement, or *GROUND*, a directionality entailing motion, or *PATH*, and a motion co-event which may be physical, i.e. *MANNER*, or causal, i.e. *CAUSE* (cf. Talmy 1985). An example of a motion event may thus be:

- (1) The rock rolled down the hill.
FIGURE MANNER PATH GROUND

The concepts of *FIGURE*, *MANNER*, *PATH*, and *GROUND* are near-systematically expressed in language when talking about motion. All languages certainly appear to encode *FIGURES*, *PATHS* and *GROUNDS*. However, these concepts are not given in the same degree of lexical codability across languages and their syntactic distribution may differ substantially. Cross-linguistic variability has been famously documented in Talmy's typological work (e.g. 1985, 1991, 2000), and in Slobin's discursive and lexical descriptions (e.g. 1996, 2000, 2003, 2004). Talmy's typology suggests that most of the world's languages follow one of two main structural patterns in encoding motion events. Languages either express the *MANNER* of motion in the main verb, and frame the *PATH* in a verb particle or verb-bound morpheme, called a satellite, as in English and other Germanic languages, among others, e.g. (1). Alternatively, languages frame the *PATH* in the main verb, and leave *MANNER* to be encoded in an optional constituent, as in French and other Romance languages, among others, e.g. (2).

- (2) Alicia a traversé le pont à vélo.
 Alicia crossed the bridge on a bike.
 FIGURE PATH GROUND MANNER
 'Alicia cycled across the bridge.'

Satellite-framed and verb-framed languages further differ in the lexical resources they employ to encode *MANNER* and *PATH*, and in the codability of those resources. The concept of *PATH* being the core schema of motion events (Talmy 1991), it is highly codable in both types of languages, though less so in verb forms in satellite-framed languages. The concept of *MANNER*, on the other hand, is highly codable in satellite-framed languages only. The difference is not only quantitative, but also qualitative, in that languages such as English afford semantic fine-graining of *MANNER* verbs, whereas Romance languages, for instance, do not.¹ Table 1 illustrates this differential codability with only a few *walk* verb examples.

These typological and lexical characteristics engender further differences at the discursive level (cf. Slobin 2004). The resulting fashions of speaking entail narratives with divergent semantic perspectives overall. English narratives, for instance, emphasise the dynamic, action- and process-oriented aspects of motion scenes; whereas Spanish and French narratives, for instance, emphasise the static, situational and resultative aspects of motion scenes (Slobin 1996). Consider, for example, the same tale reported in both language types. English may describe motion details, such as manners, very vividly, whereas Spanish may do away with these kinds of details and emphasise motion destinations or resultant states caused by motion instead.

The question of interest here is whether there might exist an ensuing correlation between this semantic distinction in the linguistic framing of motion and the non-linguistic conceptualisation of motion events by speakers of different languages.

3. Motion in language and cognition: Previous research²

Investigating linguistic relativity has been particularly popular with respect to the domain of motion, due to the importance of this domain of experience in daily human communication and conceptualisation, and also due to the importance of the linguistic differences under consideration. Based on these linguistic differences, cognitive scientists have been keen to establish whether these pervasive patterns influence speakers' cognitive representations of motion events. In particular, they have sought to correlate (i) the semantic salience of *PATH* in verb-framed constructions with the corresponding cognitive salience of the *PATH* schema in verb-framed speakers' conceptualisation of motion, and (ii) vice versa, the semantic salience of *MANNER* in satellite-framed constructions with the corresponding cognitive salience of the *MANNER* schema in satellite-framed speakers'

1. Note that only some satellite-framed languages, and not all, have a fine-grained manner-verb lexicon.

2. This review is by no means exhaustive and therefore is not fully representative of existing research.

Table 1. Examples of WALK verbs in English and French

English	French	English gloss
To march	Marcher au pas	To walk stepping
To plod	Marcher d'un pas lent	<i>To walk with a slow step</i>
To sashay	Marcher d'un pas léger	<i>To walk with a light step</i>
To saunter	Marcher d'un pas nonchalant	<i>To walk with a nonchalant step</i>
To scoot	Marcher rapidement	<i>To walk quickly</i>
To scuttle	Marcher précipitamment	<i>To walk hurriedly</i>
To shamble	Marcher en traînant les pieds	<i>To walk dragging one's feet</i>
To shuffle	Marcher en traînant les pieds	<i>To walk dragging one's feet</i>
To sidle	Marcher de côté, furtivement	<i>To walk sideways, furtively</i>
To slink	Marcher sournoisement, honteusement	<i>To walk with a mean or shameful air</i>
To slog	Marcher avec effort, d'un pas lourd, avec obstination	<i>To walk with effort, with a heavy step, with obstinacy</i>
To sneak	Marcher furtivement	<i>To walk furtively</i>
To stalk	Marcher d'un air digne ou menaçant	<i>To walk with a dignified or threatening air</i>
To stomp	Marcher d'un pas lourd, bruyant	<i>To walk with a heavy or noisy step</i>
To stride	Marcher à grands pas	<i>To walk with large steps</i>
To stroll	Marcher sans se presser, nonchalamment	<i>To walk without hurrying, nonchalantly</i>
To tiptoe	Marcher sur la pointe des pieds	<i>To walk on tiptoe</i>
To toddle	Marcher à pas hésitants	<i>To walk with hesitating steps</i>
To traipse	Marcher d'un pas traînant ou errant	<i>To walk with a dragging step or wandering aimlessly</i>
To tramp	Marcher d'un pas lourd	<i>To walk with a heavy step</i>
To trudge	Marcher péniblement	<i>To walk tediously</i>
To trundle	Marcher lourdement, bruyamment	<i>To walk with a heavy, or noisy step</i>
To wade	Marcher laborieusement dans l'eau	<i>To walk laboriously through water</i>
To waddle	Marcher comme un canard, en se dandinant	<i>To walk like a duck, lolloping</i>
To whiz	Marcher à toute vitesse	<i>To walk with great speed</i>

conceptualisation of motion (e.g. Gennari et al. 2002; Papafragou et al. 2002; Finkbeiner et al. 2002; Bohmeyer et al. 2004; Zlatev and David 2004, 2005).

In investigating this potential correlation, previous research has mainly concentrated on the cognitive abilities of categorisation and memory. One experimental format of choice has employed triadic stimuli, presented with one item as a target, and another two items as alternates. The categorisation task requires subjects to select one alternate as closer in resemblance to the target. This has, in some studies, been followed by a memory task which typically demands that the subjects recognise stimuli presented in the categorisation task. Given the experimental similarity of this type of previous research, one may well expect studies to be highly comparable in their findings. This is not the case, however.

Gennari et al. (2002) compared Spanish and English performance, using short videos of human motion events as stimuli. They report a language-independent PATH bias approximating 60%, meaning that both language groups performed similarly and that no language effect was found. However, they also used a naming-first testing condition, which yielded strong differences between the two groups, in agreement with language

patterns. Nevertheless, the authors dismiss linguistic relativity, and suggest that language influences cognition only when speakers choose to use language as a strategic problem-solving tool.

Papafragou et al. (2002) compared Greek and English performance, using static pictures and photographs depicting human motion. They report an equal distribution of *PATH* and *MANNER* scores, meaning that subjects selected either variable equally as being more cognitively salient. Besides this, cross-linguistic performance was identical in the categorisation and in the memory tasks. The authors strongly argue against relativity in their conclusions.

Finkbeiner et al. (2002) compared Japanese, Spanish and English performance, using 3-D computer animations in the shape of a ball. They report an unequal distribution of *PATH* and *MANNER* scores in favour of manner. The Japanese and Spanish language groups performed similarly, favouring *MANNER* in about 60% of their choices; whereas the English group favoured *MANNER* in 88% of choices. The authors dismiss the relevance of this significant difference, arguing that subjects used a sub-vocal linguistic memorisation technique. To support their point, they implemented the same experiment a second time but suppressed the memory element by displaying all stimuli at the same time. The manner bias found in the English group dropped dramatically. It is uncertain what may be concluded from this finding, though the authors claim it as proof of a lesser reliance on language-based strategies, and hence as a disproof of Whorfian effects. Note, however, that the other two language groups are not tested in this condition, which means that cross-linguistic data is being compared across differing experimental procedures. Overall, it remains that the data demonstrate a *MANNER* bias.

Bohnenmeyer et al. (2004) and Zlatev and David (2004, 2005) compared a high number of satellite- and verb-framed languages, using a 2-D computer animation in the shape of a smiling tomato, known as 'Tomatoman'. These research teams agree in reporting a language-independent *MANNER* bias approximating 60% in all language groups. Like the previous studies, they therefore conclude against linguistic relativity. Note, however, that Zlatev and David (e.g. 2004) do report some 'mild' effects when itemising responses as per spatial axis (i.e. vertical vs. horizontal).

In sum, previous research efforts appear to lack agreement. The one thing they agree on is that motion conceptualisation is not relative to motion encoding in language – either in categorisation or in recognition memory. However, because their findings disagree so widely, it is difficult to accept their dismissive conclusions. There are a number of problems with the above studies, however, which might explain the glaring divergences. Methodological points, for instance, shed a relative level of uncertainty over the reliability of some of the findings reported. Consider, for instance, the fact that Papafragou et al. (2002) used only eight triad items as stimuli, all of which were of a static nature (e.g. drawings) seeking to represent motion scenes which are dynamic by nature. The experiments with 'Tomatoman', on the other hand, used only 12 subjects per language group, and examined the two-dimensional motion of an artificially constructed figure, which may by no means be representative of the type of motion typically conceptualised by human subjects. Finally, Gennari et al. (2002) used an unusually high level of control in their testing instructions (see *ibid*: 62–63).

In addition to these isolated points, stimuli types differed across studies. They differed, for instance, in the type of *FIGURE* displayed as performing motion. Gennari et al. (2002) is the only study using realistic stimuli (i.e. resembling typical motion occurrences) for the representation of human motion. Finkbeiner et al. (2002), Bohnemeyer et al. (2004), and Zlatev and David (2004, 2005) obtained results pertaining to the conceptualisation of virtual, imaginary motion. If we note that the main type of motion daily conceptualised and expressed in language by speakers is that of human motion, then the relevance of these findings is questionable. Indeed, it is not quite transparent how one may relate the conceptualisation of the motion of a virtual tomato to that of human motion, or any other type of real-life, three-dimensional motion. In addition, none of the above studies contextualised motion scenes in real-life settings. The stimuli used were devoid of contextual relevance. Furthermore, most studies used a very limited number of subjects and testing formats to draw important conclusions regarding the relationship between language and cognition. In sum, it is highly possible that diverging stimuli and other methodological aspects may account for a significant portion of the variability in responses, even though the tasks were essentially the same across all research groups. These exemplars of research are therefore not always transparently comparable.

However, a more central problem pervades the above-reviewed research. The types of *MANNER* being contrasted in triadic stimuli are not fine-grained, e.g. jump vs. cycle, as opposed to hop vs. skip. Recall that the question is to assess which variable between *PATH* and *MANNER* is more cognitively salient to subjects. A triad may present [*MANNER* α + *PATH* a] in the target, [*MANNER* β + *PATH* a] in alternate (1), and [*MANNER* α + *PATH* b] in alternate (2). If the subject chooses alternate (1), s/he has categorised in terms of *PATH* similarity, thus judging *PATH* as more cognitively salient than *MANNER*; whereas if s/he chooses alternate (2), then the categorisation is in terms of *MANNER* similarity, and hence *MANNER* is judged more cognitively salient. Now, working with the cross-linguistic differences outlined above, previous research has slotted e.g. run and walk in the α and β *MANNER* types to be contrasted in one triad. In doing so, previous research has acknowledged that satellite languages would encode those *MANNER* types centrally in main verbs, whereas verb-framed languages might leave those *MANNER* types out in expression. In doing so, however, previous research has also overlooked the fact that verb languages can encode those *MANNER* types if desired. Indeed, most languages – including verb-framed ones – do have basic *MANNER* verbs such as walk, run, jump, roll, dance, swim, and so on. The linguistic difference between verb and satellite framing of motion is not merely structural, but also lexical and discursive – as mentioned above. At the level of manner distinctions, the crucial difference between those languages is that verb framing seldom affords semantic fine-graining of *MANNER* types. In other words, previous research has merely shown that Greek, Spanish, Japanese, English, and other speakers are equally able to discern running from walking, or bouncing from rolling, or jumping from tripping, and so on. Referring back to Table 1 above, for instance, it would have been more interesting to contrast α and β *MANNER* types pertaining to the same generic type of *MANNER*, e.g. types of walking, or types of running, or types of jumping, and so forth. This is precisely what Oh's (2003) research accomplished. In her doctoral work contrasting English and Korean, Oh successfully demonstrated relativistic effects in memorisation when using one

generic manner type with subtle alterations in motor control, rate or speed of performance, and the like.

Oh's research aside, most previous studies highlight one further point. Motion conceptualisation independently of language is not thoroughly understood in relativistic efforts. This important domain of experience deserves greater attention and further research. Motion has been extensively investigated in the physical, perceptual and cognitive sciences; yet, current relativistic research appears to take little account of facts of motion conceptualisation. In the present context, I propose that motion conceptualisation is not static across events (let alone across cognisers), and is fundamentally influenced by a number of variables, including *FIGURE* animacy, *MANNER* force dynamics, *PATH* telicity, and causal relations.³

4. Motion in cognition

This section presents summaries of research undertaken by Kopecka and Pourcel (2005, 2006) and Pourcel (2004, 2005) on issues of motion conceptualisation. It aims to demonstrate the relevance of fundamental properties of *FIGURES*, *PATHS*, *MANNERS*, and relational properties to understanding similarity and variability in motion conceptualisation. In so doing, it suggests that, at some given level of conceptualisation, all cognisers – regardless of their native language, cultural background, and so forth – may be similarly impressed by motion properties (see Pourcel, in press). This possibility does not preclude relativism of conceptualisation. Instead, it suggests that a number of factors enter conceptual processes, and that one such factor pertains to the very element that is conceptualised. The relativity question asks whether language may be yet another factor guiding, or shaping, conceptualisation. In addressing this latter question, one must not overlook the crucial – and possibly primary – importance of the element or event properties themselves in the conceptualisation of that element or event.


4.1 Understanding figure impact on motion conceptualisation

Kopecka and Pourcel (2005, 2006) tested native speakers of French ($N = 24$), Polish ($N = 24$) and English ($N = 21$) in order to assess the role of *FIGURE* type in motion conceptualisation. The study investigated categorisation and used triads similar to the ones mentioned in previous research. The triads comprised three types of *FIGURE*:

- a. [+animate] [-animate] *FIGURE*, i.e. the virtual tomato known as 'Tomatoman'
- b. [+animate] [+human] *FIGURE*, i.e. a real-life human being
- c. [-animate] [-human] *FIGURE*, i.e. a real-life plastic ball

The aim of the study was not to obtain relativistic effects, but to confirm which basic variables influence conceptualisation in a possibly universal fashion. For this purpose, subjects

3. The list is preliminary rather than comprehensive.



Graph 1. FIGURE type: cross-linguistic preferences for MANNER

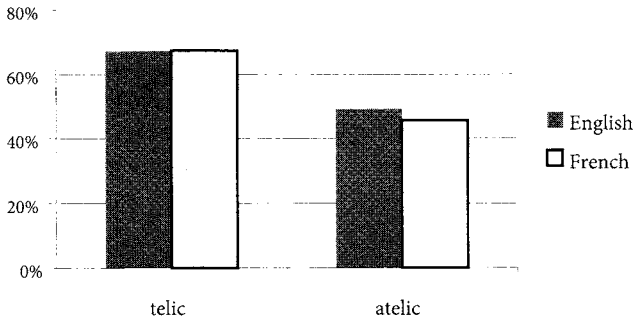
were recruited from verb-framed and satellite-framed languages and generic MANNER types only were contrasted. ‘Default’ MANNER types were used, e.g. rolling/bouncing for a ball, running/walking for the human figure, in order to minimise MANNER type interference. Likewise, only telic PATH types (i.e. with a clear endpoint) were used to avoid PATH type interference. (See next section for tests on PATH and MANNER types.) The stimuli are thus comparable in terms of PATH and MANNER types, but not in terms of FIGURE types. Finally, note that none of the motion scenes displayed caused motion. All were instances of self motion.

Results indicate a clear correlation between the relative salience of PATH and MANNER and the type of FIGURE – equally across language speakers. Indeed, as shown in Graph 1, human FIGURES correlate with a PATH bias in conceptualising motion, whereas non-human FIGURES (especially virtual ones) correlate with MANNER-focused conceptualisation.

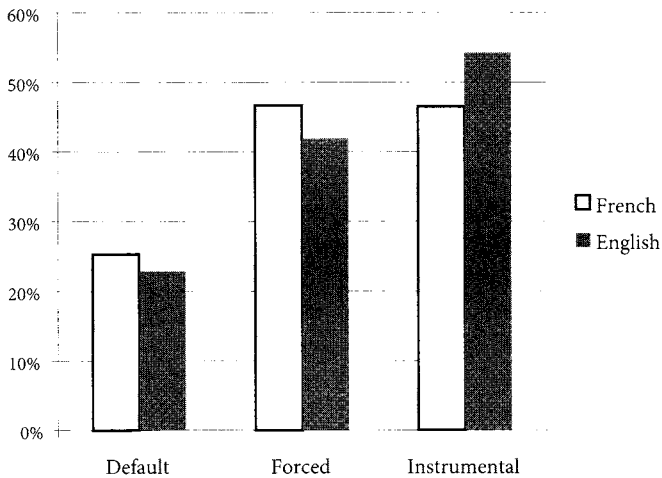
The findings reported by Kopecka and Pourcel (2005, 2006) demonstrate that previous research could not converge on their empirical results so long as they used stimuli diverging in terms of the FIGURE properties being displayed. Indeed, scores differ to significant extents relative to the FIGURE type (Mann-Whitney U-test, $p_E = 0.0004$, $p_P = 0.0002$, $p_F = 0.0002$ for Tomato-Human scores; $p_E = 0.003$, $p_P = 0.005$, $p_F = 0.003$ for Tomato-Ball scores). The present studies confirm the findings reported by most of the research reviewed previously, namely that artificial object motion triggers MANNER-salient conceptualisation of motion, whereas human motion triggers PATH-salient conceptualisation.

4.2 Understanding telicity, manner and causality impacts on motion conceptualisation

Pourcel (2004, 2005) tested 35 French speakers and 34 English speakers to assess the role of PATH, MANNER and MOTION type in motion conceptualisation. The study investigated categorisation, using 15 triads similar to the ones mentioned in previous research. The triads comprised one type only of FIGURE, namely human, hence avoiding interference from FIGURE type in associative performance. However, the triads contrasted:



Graph 2. Telicity: cross-linguistic preferences for PATH




Graph 3. MANNER type: cross-linguistic preferences for MANNER

- d. atelic and telic PATHS, e.g. along vs. across
- e. default, forced and instrumental MANNERS, e.g. walk vs. limp vs. cycle
- f. caused and self motion, e.g. X causes Y to undergo motion vs. X undergoes motion

Results indicate a significant correlation between the relative salience of PATH and MANNER and the type of PATH – equally across language speakers. Indeed, as shown in Graph 2, telic PATHS encourage a PATH bias in conceptualising motion, whereas locative, or atelic, PATHS encourage more MANNER-focused conceptualisation – though the PATH bias is merely reduced. In addition, the item score comparison reveals significant differences between telic and atelic events (Wilcoxon test $p_E = 0.001$, $p_F < 0.001$).

Results further indicate a clear correlation between the relative salience of PATH and MANNER and the type of MANNER – equally across language speakers. Indeed, as shown in Graph 3, default MANNERS correlate with a PATH bias in conceptualising motion, whereas forced and instrumental MANNERS correlate with more MANNER-focused conceptualisation – though the PATH bias is again merely reduced. The item score



Graph 4. MOTION type: cross-linguistic preferences for MANNER

comparison further reveals significant differences in conceptualisation between default and non-default events (Wilcoxon test $p_E = 0.001$, $p_F = 0.0001$ for default-forced scores; $p_E < 0.0001$, $p_F = 0.001$ for default-instrumental scores), but not between forced and instrumental items (Wilcoxon test $p_E = 0.069$, $p_F = 0.781$).

Finally, results also indicate a correlation between the relative salience of PATH and MANNER and the type of motion (i.e. self motion or caused motion) – equally across language speakers. Indeed, as shown in Graph 4, caused motion correlates with a PATH bias in conceptualising motion, whereas self motion reduces the PATH bias. The item score differences are less pronounced than in previous tests, however (Wilcoxon test $p_E = 0.067$, $p_F = 0.029$) – though this may be due to low numbers of stimuli or, possibly, participants. This suggests, nonetheless, that differences relating to causality (or lack thereof) may not cause fundamentally distinct conceptualisation of motion events.

4.3 Summary

The results reported in this section are in agreement with the results reported in the studies reviewed earlier. The present results help explain the divergences in the findings those studies report. Indeed, they show that FIGURE, PATH, MANNER, and MOTION properties are important enough to engender differential conceptualisation of basic motion events. This is true even when only one of these variables differs. In turn, this means that, to be comparable, motion event stimuli must address these differences in conceptualisation. This may be done by focusing on one type only of properties and thus narrowing conclusions as pertaining to the conceptualisation of the motion of, say, a particular FIGURE type, e.g. human motion, or object motion, or artificial motion. This may be done as well by including equivalent proportions of each type of property in the stimuli. Doing so may necessitate a cumbersome number of stimuli, which may then be limited by selecting one type of FIGURE only, or one type of motion, either self motion or caused motion, and so forth. It appears, then, that previous research has not so much failed to report convergent findings, but has more likely been dealing with stimuli divergent enough in their properties to cause conceptualisation findings to differ.

5. Motion research: Methodological considerations

The findings reported so far suggest that conceptualising motion is not a uniform process across different types of motion scenes. That is, different motion events are conceptualised differently. Conceptualisation is fundamentally influenced by the type of **FIGURE** performing the motion, by the type of **PATH** followed, the type of **MANNER** and the type of motion itself in terms of its causal properties. Given this understanding, it appears crucial to appreciate domain conceptualisation, independently of language, prior to the use and application of these domains to relativistic or other ends.

In addition, the above studies suggest that one needs to characterise more specifically than has so far been done the type of motion under focus in research. It is important to ascertain whether one is investigating human motion, object motion, animal motion, virtual motion, or other – as these appear to determine conceptualisation to extents too great to be ignored as superficial. As human motion is arguably the type of motion mostly conceptualised and expressed in language by individuals, the study to be presented here focuses specifically on human motion.

Methodologically, the studies reviewed are also valuable in highlighting a few extra critical points. First, these studies have approached the study of motion conceptualisation using constructed stimuli with minimal noise. Though the ensuing techniques are clean, they cannot help but de-contextualise the nature of ‘real’ motion, as it would be naturally conceptualised by subjects. In this sense, it may be insightful to consider testing conceptualisation of human motion as it occurs in more typical instances, that is, within real-life settings rather than out of context, as displayed in photos or short videos. I suggest that the study of motion requires more than motion events; instead, it requires a larger motion scenario, or real-life framework in which motion events are embedded and take on significance. A motion scenario would comprise internal schematic diversity, e.g. diversity of **PATHS**, fine-graining of **MANNERS**, and also dimensions external to the motion itself, e.g. agent goals, states, emotions, non-motion events, physiological senses, cultural dynamics. The following section offers an example of one such scenario.

6. Language, motion and cognition: Present research

The present study examines human motion in context. To this end, it uses a motion scenario consisting of a 4½ minute extract from the Charlie Chaplin film entitled *City Lights*. The scenario relates a suicide attempt taking place at night on a river bank with two main characters, Charlie Chaplin and a drunken gentleman. In summary, the scenario comprises:

- | | |
|----------------------------------|----------------------------------------------|
| a. three figures: | Charlie Chaplin, a millionaire, a policeman |
| b. one location: | river bank |
| c. several grounds: | river, stairs, platform, bench |
| d. several objects: | suitcase, rope, rock, flower, canes, hats |
| e. numerous MANNER types: | stumble, stroll, walk, run, limp, jump, sway |
| f. numerous PATH types: | around, down, up, into, out of, across |

- g. caused and self motion events
- h. other events besides motion: talking, crying, smelling flowers, watching
- i. psychological reality (emotions, intentions, goals, states):
fear, anger, joy, panic, surprise, tiredness
- j. cultural reality: symbols, morals, ideologies

The relativistic set of questions this study addresses are:

- i. whether French and English subjects conceptualise this same scenario differently, that is:
 - a. whether French and English subjects talk about the scenario differently
 - b. whether French and English subjects recall the scenario differently
 - c. whether French and English subjects perform inferences relating to the scenario differently
- ii. if so, how different are their conceptualisations of the scenario?
- iii. can the differences in conceptualisation be correlated with the habitual language patterns found in French and in English?


In other words, the relativistic hypothesis does not solely seek effects of language on cognition, but questions the extent of these potential effects. In doing so, the study seeks to tackle the relativistic problematic qualitatively.

For this purpose, two sample groups of French and English speakers ($N_E = 22$, $N_F = 25$) were individually observed in immediate free prose recall and late prompted recall conditions (i.e. 24 hours following visualisation). The performance of each individual was analysed for memory and inferencing.

6.1 Memory

The nature of the data was linguistic. 47 narratives were obtained from the free prose recall exercise. Narratives were analysed on an information statement-type basis, whereby a statement constitutes one type of conceptual information. This information may pertain to details of motion *PATH*, or *MANNER*, or *FIGURE*, or *EMOTION*, or *VISION*, and so on. The French narratives elicited a mean of 41.48 statements per subject, of which a mean of 21.92 statements related to motion information. The English narratives elicited a mean of 43.05 statements per speaker, including a mean of 25.23 statements relating to motion information. The quantitative performance of both groups is thus highly comparable. In addition, these motion-related statements usefully confirmed the language differences existing across French and English for expressing *PATHS* and *MANNERS* of motion (see Graph 5).

The crucial difference between French and English usage lies in the differential expression of *MANNER*. In quantitative terms, English relates *MANNER* information to a greater extent than French. Indeed, Graph 5 suggests an 18-point difference between the narratives of the two language groups. In qualitative terms, that is, in terms of semantic discrimination, the two language groups were also observed to differ. Semantic discrimination was analysed relative to token usage. A token analysis offers a count of the different lexical units for encoding *MANNER* information. The English narratives offer 88 distinct



Graph 5. Semantic representations found in cross-linguistic narratives

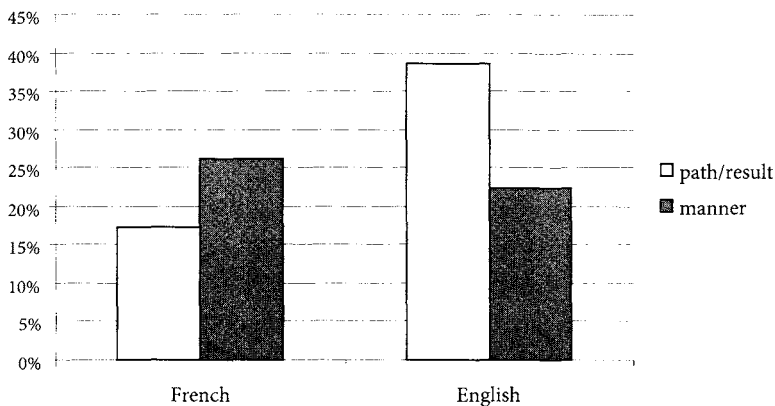
tokens encoding *MANNER*, whereas the French narratives offer 54 distinct tokens. In other words, qualitatively speaking, English usage affords a greater range of semantic values in encoding *MANNER* details.

From this template of linguistic differences concerning *MANNER* encoding, the narratives were analysed for accuracy of recall of motion variables. For this purpose, narrative statements were categorised relative to the verifiability of their truth value. In the present analysis, two main types of statements were identified relative to their truth value. Narratives in both languages presented (i) objectively descriptive statements, which offered verifiable information, that is, whose truth value could be assessed by viewing the stimulus, and (ii) subjective statements, such as ad hoc comments and inferences, which were not present in the film stimulus, but were instead reactions and ad hoc thoughts on behalf of the subject. (These are more fully discussed in the following section on inference.)

Objective statements were essentially concerned with actual descriptions of the stimulus that could be verified as correct or incorrect upon viewing. These ranged across a variety of aspects present in the stimulus, including objects e.g. (3), *GROUND*s e.g. (4), *FIGURES* e.g. (5), vision e.g. (6), text e.g. (7), *PATHS* e.g. (8), *MANNERS* e.g. (9), or scene settings e.g. (10), as shown in the following examples:

- (3) [the suitcase] has a rock and a rope in
- (4) there's water on the left of the scene
- (5) Charlie Chaplin, black hat, black moustache, black jacket, black trousers, white shirt
- (6) he looked up at what was going on
- (7) the silent movie subtitles come up saying 'tomorrow the birds will sing'
- (8) right then a policeman comes by
- (9) the drunk is always tittering on the edge
- (10) it's a black and white Charlie Chaplin film

The French narratives contained 33.32 mean descriptive statements per speaker, and the English narratives 36.77 mean descriptive statements per speaker. These objective statements constituted the vast majority of statements found in the narratives, given the



Graph 6. Error rates in the free prose recall task


descriptive nature of the task. It is equally interesting to note that 66% of the French descriptive statements related motion information ($M_F = 21.92$), as did 69% of the English descriptive statements ($M_E = 25.23$). These figures are not only comparable, they are also indicative of the high frequency of reference to motion in the present elicitations, hence confirming the relevance of this type of stimulus for investigating the domain of motion.

The aim of the free prose recall task was to monitor for similarities and differences in memory for PATH and MANNER across speakers of the two language groups. Importantly, the analysis sought to reveal whether English speakers recall MANNERS of motion better than PATHS and whether French speakers recall PATHS better than MANNERS – in line with relativistic predictions in the domain of motion. Error rates support this prediction, as shown in Graph 6, with 26% of French MANNER statements being erroneous, as compared to 17% of PATH statements, and with 38% of English PATH statements being erroneous, as compared to 22% of MANNER statements. Note, however, that these results are suggestive only of relativistic effects, as score differences are significant in the English sample only (Wilcoxon test, $p_E = 0.023$, $p_F = n.s.$).

Twenty-four hours following stimulus visualisation and the free prose recall task, subjects⁴ were asked to perform a prompted recall task. This late recall task required subjects to answer 31 questions pertaining to stimulus details, ranging from grounds to objects, causes of motion, time, figures, PATHS and MANNERS. Again, the error analysis sought to reveal whether English and French speakers displayed differential recall of PATHS and MANNERS. Error rates support this prediction, as shown in Graph 7.

Graph 7 displays different memory scores for PATHS and MANNERS by the French and the English sample groups. English subjects display better recall of MANNERS than French subjects, and French subjects display better recall of PATHS. These results concord with relativistic predictions. However, these preliminary results are suggestive only (Mann-Whitney U-test, $p_{PATH} = 0.05$, $p_{MANNER} = n.s.$).

4. Note that a greater number of subjects completed this task ($N_F = 33$, $N_E = 29$).



Graph 7. Error rates in the late recall task

6.2 Summary

The recall analyses have demonstrated that French and English speakers talk differently about motion scenarios, and recall these scenarios differently too. Indeed, the French test sample displays better recall of PATHS, whereas the English sample displays better recall of MANNERS of motion. These differences correlate with the conceptual representations receiving linguistic emphasis in each language. This correlation is therefore suggestive of relativistic effects in memory.

An interesting question to contemplate from this suggestive evidence concerns the kinds of entailments that may be drawn from the differences thus far observed. Upon consideration of cross-linguistic data, Slobin (1996: 84) has suggested that “English speakers tend to assert actions, implying results, whereas Spanish speakers assert results, implying actions.” Given that we now have preliminary evidence that such may be the case at the cognitive level too, I suggest asking two further questions:

- a. Do English speakers conceptualise motion in more dynamic and processual terms (given the dynamic and temporal dimensions of MANNERS of motion)?
- b. Do French speakers conceptualise motion in more configurational and resultative terms (given the relative lack of attention to MANNERS and the end-oriented dimension of PATHS, especially telic ones)?

In other words, one may ponder the greater significance of PATH and MANNER schemas to larger and maybe more systematic patterns of thought. The suggestion made here – also following from Slobin’s comment – is that attention to MANNER may correspond to attention to dynamic aspects of motion scenes, and that attention to PATHS may correspond to attention to possibly more static or configurational aspects of motion. Indeed, MANNER seems to relate to the motion itself, as well as to the motor capacities of the moving figure, whereas PATH appears to relate to specific points of a motion trajectory relative to given grounds, as well as to the goal and purpose of the moving figure. The following speculation is thus offered, according to which MANNERS call attention to the

dynamic and physical/motor properties of motion, whereas PATHS call attention to the configurational and mental/intentional properties of motion – at least insofar as human motion is concerned. We may call this speculation the 'human motion salience hypothesis'. Dynamic and motor properties differ in important ways from configurational and intentional properties of events. In fact, these highlight very distinct perspectives. The question then becomes much broader and asks whether French and English speakers have grown pre-disposed and accustomed to adopting the one or the other perspective when processing thoughts relating to motion, and maybe even to other events.


6.3 Inference

To examine these broader possibilities in the domain of motion, I suggest analysing the subjectivity present in the Charlie Chaplin narratives. As mentioned in the previous section, the narratives yielded two types of statements: (i) objective descriptions and (ii) subjective comments. As opposed to objective descriptions, subjective statements could not be verified as either true/false, or accurate/inaccurate, and they could therefore not be subjected to recall analysis. The subjectivity present in the narratives proves highly interesting, nonetheless, as these statements were largely inferential in nature. In addition, subjectivity proved to be present to a significant extent in the narratives of both sample groups. 20% of all French statements were subjective ($M_F = 8.16$), as were 15% of all English statements ($M_E = 6.27$).

These reported means suggest some differences between the two rhetorical styles of the language groups under study, with French speakers employing subjectivity and inferentiality to a slightly greater extent than English speakers. Note, however, that cross-linguistic performance remains comparable. The present analysis focuses on those inferential subjective statements. The type of inferences present in the narratives ranged across likely actions e.g. (11), likely results of actions e.g. (12), likely intentions or goals e.g. (13), likely thoughts e.g. (14), likely psychological states e.g. (15), likely communicative events e.g. (16), or likely material details such as ground peculiarities e.g. (17).

- (11) ...as if he's going to hang himself.
- (12) ...basically the guy saves Charlie.
- (13) ...the man wants to kill himself.
- (14) ...Charlie manages to put two and two together.
- (15) ...he's very serious and sad.
- (16) ...sort of says 'good night, how are you?'
- (17) ...the water must be quite shallow because they don't drown.

The kind of information, as illustrated above, that is of an inferential nature, is particularly interesting to linguistic relativity, because it offers rich glimpses at habitual ways of thinking about a given scene. These habitual ways of thinking are very individual. But according to the linguistic relativity hypothesis, they should also be partly influenced by the fashions of speaking in particular languages (Whorf 1956; Lucy 1996). Indeed, subjective statements may be dispensed of when narrating a tale, especially when the narrating task



Graph 8. Cross-linguistic distribution of inference types in the narratives


focuses on description.⁵ Hence, when they are employed, subjective statements encode information that is perceived by the narrator as highly relevant and salient. If, indeed, habitual ways of thinking and of inferring are influenced by habitual ways of speaking, then we may expect English and French speakers to make different types of inferences to different extents.

As mentioned earlier, *PATHS* and *MANNERS* of motion may be said to relate to configurational, cognitive and intentional properties, and to active, dynamic and physical properties, respectively. In this respect, all types of inferences outlined above are of special interest to the present analysis.⁶ The relativist hypothesis would, in turn, predict that the semantic emphasis on *MANNERS* in English should cause (a) the manner schema to be more cognitively salient to English speakers, and hence better recalled in memory (which has been the object of the previous section), and likewise (b) the concepts associated with that schema to be more cognitively salient to English speakers, and hence more prone to inferential reference in subjective performance. In other words, English speakers are expected to make more inferences bearing on likely actions, for instance. Conversely, French speakers are expected to draw inferences of a cognitive, or intentional, type more readily than English speakers. Categorising inferential statements relative to these distinctions reveals a differential distribution of information, as shown in Graph 8.

Graph 8 demonstrates that both groups perform all kinds of inferences in discourse. However, interestingly, each group differs from slightly to greatly in the type of information inferred. These differences are noteworthy when contrasting dynamic processes, such as actions, and agent-centred properties pertaining to mental aspects of motion processing, such as intentions – in the case of actions and intentions, score differenc-

5. Note that only one subject offered a narrative devoid of subjectivity. The subject was a native English speaker.

6. Note that too few inferences on likely grounds were present for an analysis of configurational representations.



Graph 9. Cross-linguistic distribution of dynamic versus cognitive inference types

es are significant across language groups (Mann-Whitney U-test, $p_{\text{ACTIONS}} = 0.0002$, $p_{\text{INTENTIONS}} = 0.036$). Graph 9 conflates the two types of information in a bar chart for a clearer illustration of this point.

These results suggest considerable differences across the two language groups concerning their habitual ways of conceptualising and inferring motion dimensions (Mann-Whitney test, $p < 0.05$ for cross-linguistic scores for both dynamic processes and for cognitive states; Wilcoxon test, $p_E = 0.046$, $p_F < 0.005$ for item scores within languages). Overall, the English group infers information of a more dynamic and processual nature than does the French group, whose inferences relate more readily to likely agentive mental states.

6.4 Summary

Inferential information represents optional elements for the purpose of narrative reporting. The information choices operated by each language group show that English speakers encode action-based aspects of motion more frequently than French speakers in their inferences. These choices also display important discrepancies between the two groups with respect to information of an interpretive, intentional, and psychological/emotional type. In other words, French speakers elaborate to a greater extent on the cognitive/mental aspects involved in motion, whereas English speakers elaborate to a greater extent on the dynamic and physical reality of motion.

The findings reported in this study indicate that the two languages employ different narrative styles indeed. Whether objective or subjective, French and English tokens and statements differed on specific types of information. English narratives conveyed more information relating to *MANNERS* of motion and to action-based inferences, hence emphasising those aspects of motion which are physical, dynamic, and overt (i.e. readily observable). French narratives, on the other hand, devoted less emphasis to dynamic motion details and, instead, focused to a greater extent on interpretations of figures, intentions and emotional states relating to motion. In other words, French narratives devoted greater emphasis to motivational, mental, and covert aspects of motion. In this sense, we may suggest that English rhetoric is more concerned with describing the physical reality of motion, whereas French rhetoric appears to be more concerned with interpreting the cognitive reality behind motion. Thus, as predicted and widely documented by Slobin's

extensive narrative research, Germanic and Romance languages differ in their narrative styles for framing motion semantics. These divergences are not purely lexical, but emerge through the informational choice and selection operated in discourse production, and they contribute to defining specific narrative, or rhetorical, styles of expression in each language.

7. Conclusion

To conclude, it is hoped that this paper has made a number of contributions to the empirical study of linguistic relativity in the domain of motion. These contributions have been both methodological and empirical.

Methodologically, this paper has suggested that a choice of motion stimulus that takes context into account might better inform realistic conceptualisation than isolated and artificial stimuli. To illustrate this point, this paper has offered an exemplar of a motion scenario, itself consisting of several motion events, yet crucially embedding these motion events within a meaningful environment including other events, agent motivations, action consequences, psychological states, and cultural significance.

A review of past research has led this paper to stress the importance of domain understanding, independently of language, prior to relativistic applications. In the present study, this has been illustrated via an exploration of the potential universals in motion conceptualisation. Several fundamental motion components were identified as impacting on conceptualisation, including *FIGURE* type, *PATH* telicity, *MANNER* force dynamics, and motion causality.

Taking this understanding into account, the research reported in this paper sought to implement a more informed methodological approach to relativity testing than has so far been used in motion research. To this end, it has used a motion scenario to examine motion conceptualisation in memory (immediate and late recall) and in inferencing. The results reported suggest important differences across the experimental groups in linguistic and cognitive terms. To the set of relativistic questions asked earlier in (i)–(iii), we may thus answer as follows:

- i. French and English subjects conceptualise the same motion scenario differently, that is:
 - a. French and English subjects talk about the scenario differently.
 - b. French and English subjects recall the scenario differently.
 - c. French and English subjects perform inferences relating to the scenario differently.
- ii. Significant differences are apparent in their conceptualisations of the scenario, i.e.
 - a. French subjects recall *PATHS* more accurately and draw more inferences concerning cognitive states.
 - b. English subjects recall *MANNERS* more accurately and draw more inferences concerning dynamic processes.

- iii. The cross-linguistic differences in conceptualisation correlate with the habitual language patterns found in French, which emphasise PATHS in motion event expression, and those found in English, which emphasise attention to both PATHS and MANNERS in motion event expression.

In other words, the findings reported in this study suggest differing cognitive styles in motion conceptualisation. These cognitive styles parallel linguistic differences very closely, and may therefore constitute correlational evidence for linguistic relativity in the domain of motion, across the French and English speaking communities.

A final contribution offered in this paper has been to launch a discussion relating specific schemas, such as PATH and MANNER, to grander patterns of thought. The argument may need further definitional precision at this point, yet it is important to link local findings to broader issues of conceptualisation. Indeed, relativity studies often content themselves with proving local effects without suggesting the pertinence of these effects to human daily lives. Experimental data should ideally be used as a stepping-stone for a deeper understanding of conceptual representations. Such a step is crucial to showing the relevance of laboratory findings to human lives and to bridging the gap between individual and collective cognition. Indeed, a consideration of broader patterns of thought may lead us towards methodologies for tackling the 'untouchable' notion of worldview so central to the Whorfian paradigm.

References

- Aske, J. 1989. Path predicates in English and Spanish: A closer look. *Proceedings of the Fifteenth Annual Meeting of the Berkeley Linguistics Society*: 1–14.
- Berlin, B. & Kay, P. 1969. *Basic Color Terms*. Berkeley: University of California Press.
- Bohnenmeyer, J., Eisenbeiss, S. & Narasimhan, B. 2004. Manner and path in non-linguistic cognition. Conference presentation, Language, Culture and Mind, Portsmouth.
- Danziger, E. 2001. *Relatively Speaking: Language, Thought, and Kinship among the Mopan Maya*. Oxford: Oxford University Press.
- Finkbeiner, M., Nicol, J., Greth, D. & Nakamura, K. 2002. The role of language in memory for actions. *Journal of Psycholinguistics Research* 31 (5): 447–457.
- Gennari, S. P., Sloman, S. A., Malt, B. C. & Fitch, W. T. 2002. Motion events in language and cognition. *Cognition* 83: 49–79.
- Kopecka, A. & Pourcel, S. 2005. Figuring out figures' role in motion conceptualisation. Conference paper presented at the 9th International Cognitive Linguistics Conference. Seoul, Korea.
- Kopecka, A. & Pourcel, S. 2006. Understanding FIGURES in conceptualising motion. Conference paper presented at the 2nd Biennial Conference on Cognitive Science. St Petersburg, Russia.
- Levinson, S. C. 2003. *Space in Language and Cognition*. Cambridge: Cambridge University Press.
- Lucy, J. A. 1992. *Language Diversity and Thought*. Cambridge: Cambridge University Press.
- Lucy, J. A. 1996. The scope of linguistic relativity: an analysis and review of empirical research. In *Rethinking Linguistic Relativity*, J. J. Gumperz & S. C. Levinson (eds), 37–69. Cambridge: Cambridge University Press.
- Lucy, J. A. 1997. Linguistic relativity. *Annual Review of Anthropology* 26: 291–312.
- Oh, K.-J. 2003. *Habitual Patterns of Language Use and Thinking for Speaking: A Whorfian Effect on Motion Events*. PhD Dissertation, University of California, Berkeley.

- Papafragou, A., Massey, C. & Gleitman, L. 2002. Shake, rattle, 'n' roll: the representation of motion in language and cognition. *Cognition* 84: 189–219.
- Pourcel, S. 2004. Motion in language and cognition. In *Linguagem, Cultura e Cognição: Estudos de Linguística Cognitiva*, A. Soares da Silva, A. Torres & M. Gonçalves (eds), Volume 2: 75–91. Coimbra: Almedina.
- Pourcel, S. 2005. What makes path of motion salient? *Proceedings of the Thirtieth Annual Meeting of the Berkeley Linguistics Society*: 505–516.
- Pourcel, S. in press. Motion: A conceptual typology. In *Language, Cognition and Space: The State of the Art and New Directions*, V. Evans & P. Chilton (eds). London: Equinox Publishers.
- Slobin, D. I. 1996. From 'thought and language' to 'thinking for speaking'. In *Rethinking Linguistic Relativity*, J. J. Gumperz & S. C. Levinson (eds), 70–96. Cambridge: Cambridge University Press.
- Slobin, D. I. 2000. Verbalised events: A dynamic approach to linguistic relativity and determinism. In *Evidence for Linguistic Relativity*, S. Niemeier & R. Dirven (eds), 107–138. Amsterdam: John Benjamins.
- Slobin, D. I. 2003. Language and thought online: Cognitive consequences of linguistic relativity. In *Language in Mind: Advances in the Investigation of Language and Thought*, D. Gentner & S. Goldin-Meadow (eds), 157–191. Cambridge, MA: MIT Press.
- Slobin, D. I. 2004. The many ways to search for a frog: Linguistic typology and the expression of motion events. In *Relating Events in Narrative: Typological and Contextual Perspectives*, S. Strömquist & L. Verhoeven (eds), Volume 2: 219–257. Mahwah, NJ: LEA Publishers.
- Talmy, L. 1985. Lexicalisation patterns: Semantic structure in lexical forms. In *Language Typology and Syntactic Description*, T. Shopen (ed.), Volume 3: 57–149. Cambridge: Cambridge University Press.
- Talmy, L. 1991. Path to realisation: A typology of event conflation. *Proceedings of the Seventeenth Annual Meeting of the Berkeley Linguistics Society*: 480–520.
- Talmy, L. 2000. *Toward a Cognitive Semantics: Typology and Process in Concept Structuring*. Volume 2. Cambridge, MA: MIT Press.
- Whorf, B. L. 1956. *Language, Thought, and Reality*. Cambridge, MA: MIT Press.
- Zlatev, J. & David, C. 2004. Do Swedes and Frenchmen view motion differently? Conference paper presented at Language, Culture and Mind. Portsmouth, UK.
- Zlatev, J. & David, C. 2005. Motion event typology and categorisation. Conference paper presented at the 9th International Cognitive Linguistics Conference. Seoul, Korea.
- Zlatev, J. & Yangklang, P. 2004. A third way to travel: The place of Thai in motion–event typology. In *Relating Events in Narrative: Typological and Contextual Perspectives*, S. Strömquist & L. Verhoeven (eds), 159–190. Mahwah, NJ: LEA Publishers.

PART V

Extensions and applications of cognitive linguistics

Toward a social cognitive linguistics

William Croft

1. Introduction

Cognitive linguistics aspires to be an approach to language (see e.g. Evans and Green 2006: 3), in particular an alternative to formal (Chomskyan) theories of grammar and formal theories of semantics. The basic principles of cognitive linguistics are often formulated as rejections of basic principles of formal syntax and semantics. In their place, the principles of cognitive linguistics offer an alternative which is more plausible and fruitful for understanding the nature of language (at least in the minds of its practitioners and sympathizers). Yet cognitive linguistics is in danger of construing itself too narrowly as an approach to *language*, in the same ways that formal syntactic and semantic theories have been criticized as too narrow. This is not to say that the foundations of cognitive linguistics are invalid. They do offer a model of linguistic cognition that has greater potential than the formal alternatives, in my opinion at least. But they are incomplete. In particular, as my title implies, they are too solipsistic, that is, too much ‘inside the head’. In order to be successful, cognitive linguistics must go ‘outside the head’ and incorporate a social–interactional perspective on the nature of language.

Cognitive linguistics can do so by incorporating certain foundational work in pragmatics and sociolinguistics. Integrating these two perspectives, the cognitive and the social, would be an important step forward in providing a genuine approach to the whole of language. There is a long-standing separation of the social and psychological dimensions in the study of human behavior, language included. This gap must be bridged in order to achieve progress in understanding the nature of language. There are a few important antecedents in linguistics, particularly the work of Talmy Givón and Wallace Chafe. In my chapter here, I will largely draw on the work of linguistic and philosophical pragmatics, as well as some work in sociolinguistics. I will also draw heavily on the interpretation of pragmatic research by the psycholinguist Herbert H. Clark, who has argued for a comprehensively and consistently social cognitive perspective on language (Clark 1992, 1996, 1999). I also take inspiration from the work of the psychologist Michael Tomasello (e.g. Tomasello 1999, 2003, 2008). In the last part of this chapter, I will demonstrate the fruitfulness of this approach in addressing traditional cognitive linguistic questions, in particular the nature of construal and its relation to grammar.

2. Basic principles of cognitive linguistics and their shortcomings

Perhaps the most fundamental distinguishing characteristic of cognitive linguistics as a research paradigm is the hypothesis that *grammatical structures and processes in the mind are instances of general cognitive abilities*. In other words, language is not an autonomous cognitive faculty (Croft and Cruse 2004: 1). The latter, negative formulation from Croft and Cruse is explicitly a rejection of generative grammar's fundamental philosophical premise. The positive formulation (see Croft and Cruse 2004: 2–3) has been extremely fruitful, allowing cognitive linguists to draw on important research in cognitive psychology and Gestalt psychology on categorization, prototypes, memory, attention and so on in order to illuminate linguistic phenomena.

A second principle which has guided much work in cognitive linguistics is that *grammar is symbolic, and thus meaning is an essential part of grammar*. The first part of this principle underlies construction grammar as a model of grammatical organization (Fillmore, Kay and O'Connor 1988; Goldberg 1995, 2006; Croft 2001; Croft and Cruse 2004, part III). The second part explains the emphasis on semantics in the majority of research in cognitive linguistics. In both respects, cognitive linguistics is again reacting to generative grammar. Generative grammar has emphasized research on syntax at the expense of semantics. Narrowly construed, generative grammar is a theory of syntax; more broadly, syntax interfaces to 'lexical conceptual structure', which appears to be beyond the purview of generative grammar. Some semanticists allied to the generative approach to syntax work in the formal semantic tradition, while others, such as Ray Jackendoff, are closer to cognitive linguistics than to formal semantics. Generative grammarians however do include a semantic component in their grammatical models. The real difference between generative grammar and construction grammar is that in construction grammar, syntax and semantics are not partitioned into autonomous components, but united as a 'structured inventory of conventional symbolic units' (Langacker 1987: 57). (Even here, one offshoot of generative grammar, Head-driven Phrase Structure Grammar, has taken a symbolic approach to grammatical organization.) A further corollary to this principle is that even grammatical constructions have meanings.

The third and fourth basic principles of cognitive linguistics focus on meaning, and differentiate cognitive semantics from formal, logic-based, truth-conditional semantics. The third principle is that *meaning is encyclopedic*. This principle implies that one cannot separate a subset of semantic features or predicates as constituting "the meaning" of a word or a construction. Instead, all that the speaker knows about the real world experience denoted by the word or construction plays a role (however small) in its meaning. One important way in which encyclopedic knowledge plays a role in word meaning is described by the frame semantic model (Fillmore 1982, 1985, *inter alia*; for a textbook survey, see Croft and Cruse 2004, Chapter 2). Frame semantics hypothesizes that the meaning of a word includes its background presuppositions or semantic 'frame' and cannot be understood apart from its frame. Many aspects of the semantic frame for a word are not truth-conditional, and therefore truth-conditional semantics is an incomplete theory of linguistic meaning, according to cognitive semanticists.

The fourth and final principle is that *meaning involves conceptualization (construal)*. That is, linguistic meaning includes our perspective on a particular state of affairs. Again, a truth-conditional characterization of the meaning of a linguistic expression is insufficient because it usually makes reference to only the state of affairs itself in defining truth conditions, not the conceptualization of the state of affairs by the speaker. This principle of cognitive semantics therefore also distinguishes it from formal semantics. The principle of conceptualization is interconnected with the other three principles. The framing of an experience through the choice of a lexical item is a matter of construal. The construal operations proposed by cognitive linguists are analyzable as instances of general cognitive processes (Croft and Cruse 2004, Chapter 3; see also §3.4). Finally, the meanings of grammatical constructions and grammatical elements, often treated by generative grammar as lacking meaning, have been analyzed by cognitive linguists as having as their function the imposition of a particular construal on a state of affairs.

These four principles capture much of what all cognitive linguists would agree upon. All of them focus on language as a cognitive ability, in terms of how language is represented in the mind, what is represented, and what cognitive processes are involved in the production and comprehension of language. This is after all why this approach is called “cognitive linguistics”: it is a model of what language is as a cognitive ability.

But language is not just a cognitive ability, a constellation of mental structures and processes. If it were just something mental, we would not need to speak. We might not even need language at all; we would just have our mental representations of meaning. Yet that is patently not the case. We all have language and use it. The reason, of course, is that language is a central feature of human social interaction. But this means that language cannot be fully understood outside of that fact. An approach to language such as cognitive linguistics, as it is presently constituted, therefore cannot provide us with a complete understanding of the nature of language. Cognitive linguistics must reach out and embed itself in a more general social–interactional model of language (Croft and Cruse 2004 concludes with this observation). Or to construe the issue differently, we must bring together the cognitive and the social dimensions of language, in a way that has hardly been done before by either cognitively or socially/ functionally oriented linguists (but see Sinha 1999 and Verhagen 2005 as well as the scholars mentioned in §1). I will propose a way to do so by reformulating the four basic principles of cognitive linguistics presented in this section.

3. Rethinking cognitive linguistics socially

3.1 From general cognitive abilities to general social cognitive abilities

The first basic principle of cognitive linguistics outlined in the preceding section is that *grammatical structures and processes in the mind are instances of general cognitive abilities*. The general cognitive abilities that are appealed to by cognitive linguists can be found in textbooks on cognitive psychology under the headings of perception, memory and categorization. Perceptual principles are involved in semantic representation and

the comprehension of utterances. Memory is involved in the organization of grammatical knowledge and the production of utterances. Categorization is involved in all of the above, since it plays a central and pervasive role in human cognition.

All of these are of course essential for understanding the nature of language. But in addition to these cognitive abilities, we must also recognize that *grammatical structures and processes in the mind are instances of general social cognitive abilities*. The most important of these social cognitive abilities are joint action, coordination and convention.

Language is a joint action, like many other human actions. All too often it is thought of as something individual: the speaker formulates a thought and produces an utterance. Or a listener hears an utterance and activates a meaning as a consequence. While these individual actions are certainly a part of language, they do not capture what is special about it, namely that speaker and hearer are together engaging in a joint action. A speaker speaks with an audience in mind; we think of someone speaking without an audience as a disordered individual (or nowadays, probably speaking to someone on their cellphone). A listener more obviously cannot listen without a speaker, but it must not be forgotten that the listener is not just activating a meaning of the utterance but is attempting to understand what the speaker intends to communicate. Linguistic communication is an instance of the more general human social cognitive ability of engaging in joint action.

The description of joint action presented here follows the philosopher Michael Bratman's definition of shared cooperative activity (Bratman 1992; see also Bratman 1993, 1997). It is close to the definition of joint action found in Clark (1996:60–62) but adds certain conditions that Clark's theory lacks or is inexplicit about (which is partly because Bratman's definition may define a narrower class of actions than Clark's). Loosely, what makes a joint action joint is that it is more than just the sum of individual actions performed by separate persons; in particular each individual involved must take into consideration the other individual's beliefs, intentions and actions in a way that can be described as cooperative. We can illustrate the joint action with two examples: two persons carrying a sideboard into a dining room, and two persons performing a sonata for cello and piano. Bratman argues that a joint action such as these two must have the following specific features:

Each of us intends to perform the joint action. As the pianist, I intend that I *and* my cellist should perform the cello sonata together. This is of course more than my intending just to perform the piano part (see below). Likewise, I intend that *we* carry the sideboard into the dining room; my individual action doesn't make sense by itself.

Each of us intends to do it in accordance with and because of each of our meshing subplans. Bratman's term 'meshing subplans' makes reference to the individual actions and subactions carried out by each individual. I intend that we perform the cello sonata by performing the piano part, but the way that I play the piano part – the tempo, emotional interpretation, etc. – is intended to mesh with my cellist's performance of the cello part. In turn, my performance of each movement of the piano part, each measure of each movement, and each beat of each measure must mesh with the corresponding movement, measure and beat of the cello part. That is, the subplans can be sequenced and nested, and they must all mesh with the other individual's subplans. Likewise, I intend that we should move

the sideboard by lifting a certain end of it and moving in a certain direction, but that action must mesh with the corresponding actions of my colleague.

Neither of us is coercing the other. Certain situations involve coercion without any intentions on the part of the other: I kidnap you and throw you into the trunk of the car, and ‘we’ go to New York (Bratman 1992: 333). This is clearly not a joint action. But I could alternatively point a gun to your head and force you to go with me to New York, or otherwise coerce you into doing something with me (ibid.: 334–335). Once that matter is settled, you will then mesh your subplans with mine, like it or not. Bratman includes this condition in order to exclude such examples from shared cooperative activity.

Each of us has a commitment to mutual support. That is, each of us has a commitment to help each other carry out their subplans. If my cellist misses a note, or falters in some way, I can attempt to fill in, keep up the tempo, or any one of a number of things I can do to ensure that the joint action is carried out successfully to the best of our abilities. Likewise, if my colleague carrying the sideboard slips, I will stop or otherwise do what I can to prevent the sideboard from being damaged so that when he is ready again, we can successfully carry out the action. Tomasello (2008) suggests that this particular condition, helpfulness, is perhaps one of the significant evolutionary steps to cooperative activity among humans.

This is common ground (shared knowledge, belief, suppositions) between the individuals. That is, the intentions described above must be common ground among the individuals performing the joint action. It must be common ground between us that I intend to carry out the action of jointly carrying the sideboard, and that I will mesh my individual actions with yours, and the same for you to me. Otherwise you and I will not succeed in the task, and each of us may not have any incentive to try to carry it out. Clark argues, following Lewis (1969), that common ground should be defined in terms of a shared basis. Common ground will be discussed in more detail in §3.3; the crucial fact for now is that common ground exists and that it is more than just individual knowledge or belief about someone else’s mental states.

There is mutual responsiveness in action. That is, in executing the joint action we will *coordinate* our individual actions in order to ensure that they mesh with each other in execution and hence the joint action will be successfully carried out (to the best of our abilities). Coordination is essential in carrying out joint actions successfully and I will discuss it in greater detail below.

All of these features are, or depend on, general social cognitive abilities. The ability to conceptualize a joint action, not just individual actions, is a social cognitive ability. The remaining features are also social abilities: the ability to mesh individual activities, the capacity for common ground, the ability for mutual support with the goal of the joint action, and the ability to coordinate actions (be mutually responsive).

All of these abilities involve not just one individual but two (or more) individuals interacting in a way that presupposes among other things the ability to recognize a conspecific as having intentions and other mental states like one’s own (Tomasello 1999; Étienne Wenger, whose work is discussed in §3.3, calls this ‘mutuality’ [Wenger 1998: 56]). More recently, Tomasello (2008) argues on the basis of experiments and observations of human infants and nonhuman primates, that cooperation – the ability to engage in joint actions

of the type described by Bratman – is the root of other social cognitive abilities and is a special, possibly unique characteristic of human beings. In fact, we can probably safely say that it is the human ability to carry out cooperative joint actions that has led to the creation of human society and the vast array of cultural and technological achievements that are so obviously manifest and define much of our way of being in the world, far different from that of other species.

Language is itself a joint action, as I mentioned before. The joint action is communication, which is often described as the “function” of language. (I will qualify this statement shortly.) When speaker and hearer converse, they intend to perform the joint action of communication. Each intends to perform it in accordance with and because of their meshing subplans. These subplans – individual actions – are described in greater detail in §3.4, but for now we can describe them as uttering something (the speaker) and listening and understanding it (the hearer). These intentions are common ground: why should I talk to you if there is no shared knowledge that you are listening? Speaker and hearer are not coercing each other (under normal circumstances). Each of us has a commitment to mutual support: if you don’t understand me, or can’t hear me, you ask me to clarify or repeat, and I do so. Finally, there is coordination in our actions: we monitor each other’s actions and attention in order to ensure that our subplans are meshing properly, and that the joint action is being carried out successfully.

Coordination plays a particularly important role in the successful performance of joint actions. Without it, our intentions to make our subplans mesh will fail. For example, in carrying the sideboard, we will inevitably drop it if we do not coordinate our individual actions so that they mesh properly. Human beings have a number of coordination devices that allow joint actions to happen. One of the most important ones is the ability to establish *joint attention*. In moving the sideboard into the room, we can jointly attend to the position of the sideboard, the location of the door and the doorjamb that we don’t want to bump into, and so on. We can ensure that joint attention is established by pointing (in the case of moving the sideboard, with our nose or chin), or by looking at each other (as frequently happens with musicians playing together), to ensure joint attention for coming in together at certain points in the music. The ability to establish joint attention is also something apparently unique to humans, or at least humans have the greatest facility in establishing joint attention (Tomasello 1999, Ch. 3). Be that as it may, joint attention is certainly a social cognitive ability that is a prerequisite for language – at the very least, as I noted above, speaker and hearer must jointly attend to the utterance being produced.

Nevertheless, joint attention limits one in the sort of joint actions that can be accomplished. It allows coordination of current actions in the here and now, and is essential for doing so, but it hardly allows for planning or more complex coordination. For example, in moving the sideboard, coordination may require a more precise description of the path of motion (to avoid the doorjamb) and the destination of the sideboard. Or in playing the cello sonata, we might want to coordinate certain changes in tempo and dynamics as we go through a particular movement. A far more powerful coordination device that humans have the ability to use is *communication*. I can communicate to you to watch for the doorjamb which you can’t see because it’s behind you and you’re keeping an eye on the sideboard; or I can communicate to you a plan to slow the tempo and increase the dynam-

ics at a major cadence in the cello sonata. In fact, communication is essentially a coordination device for joint actions (Clark 1999). Linguists often think of communication as the conveyance of information. But even in the conveyance of information, there is a joint action: namely, the creation of shared information (common ground). Communication, linguistic or otherwise, is just the coordination device for that purpose and many others (Clark 1999).

The power of communication as a coordination device for joint actions is that communication allows an individual to recognize the intentions of the other individual and make those intentions part of the pair's common ground. But communication is itself a joint action, even as it is functioning as the coordination device for other joint actions that human beings want to engage in. Here we run into a fundamental fact about our social existence. We cannot read each other's minds (cf. Croft 2000:95). A moment's reflection should allow the reader to recall times in which this is not such a bad thing. Nevertheless, this fundamental social cognitive fact means that successful communication is a major challenge. Communication itself needs a coordination device. The most powerful coordination device for communication is convention, specifically linguistic convention. This is a major part (though not all) of what language is.

Hence another social cognitive ability necessary for understanding the nature of language is convention. The definition of convention that I will use is originally derived from the philosopher David Lewis (1969), but it has been reformulated by Herbert Clark in a way that highlights how it is joint in nature. I present Clark's reformulation, with further specifications given by Lewis, using two examples of convention: shaking right hands as a convention for greeting someone, and a linguistic convention such as the use of the string of sounds *butterfly* to refer to a particular insect. The combined Lewis and Clark definition of convention is given below (see also Croft 2000:98):

A regularity in behavior... The behavior is choosing to shake right hands, or producing the string of sounds notated in English as *butterfly*. This is a regular behavior; that is, it is one that is used on repeated occasions.

...that is partly arbitrary. Lewis defines arbitrariness as the situation in which other regularities in behavior would be approximately equally preferable by almost everyone in the relevant community (see below), e.g. shaking with the left hand instead of the right, or using the string of sounds *mariposa* instead.

...that is common ground in a community. As with joint actions, the use of linguistic convention must be common ground (shared knowledge) among the people who use it. This group of people are the relevant community, e.g. the cultural group that shakes right hands or the speech community that uses *butterfly*.

...as a coordination device. Convention is specifically a coordination device, not unlike joint attention. But convention specifically requires for its success that members of the community conform to it to a great degree. As Lewis puts it more precisely (and carefully hedged): almost everyone in the community conforms to it; almost everyone expects almost everyone else to conform to it; almost everyone would prefer any additional member of the community to conform to it if almost everyone in the community already conforms to it; and almost everyone would prefer any new member of the community to conform to another regularity if almost everyone in the community were already conforming to it.

...for a recurrent coordination problem. The recurrent coordination problem is to greet someone, in the case of shaking hands, or to communicate the meaning of referring to individuals belonging to a particular category of insect, in the case of *butterfly*. The crucial condition for the evolution of conventions is that the coordination problem is recurrent: we frequently wish to greet someone (a joint action), and we frequently wish to refer to individual entities such as butterflies (communication; also a joint action). The fact that the coordination problem is recurrent makes it sensible for a community to settle upon a convention to solve it.

The ability to evolve conventions for communication is a social cognitive ability essential for language. Indeed, for many people language *is* the set of linguistic conventions of a speech community. Although this is too restrictive a view, convention is certainly the most central and in many respects the most distinctive characteristic of human linguistic behavior. Nevertheless, convention and also coordination are general social cognitive abilities; they are not specific to language, and language is not a special mental module by virtue of being a conventional coordination device. I have already given one example of a nonlinguistic convention, namely greeting someone by shaking right hands. Other greeting conventions exist as well, such as kissing on the cheek (and the number of times also varies conventionally from culture to culture). Other conventions for coordination of other kinds of joint actions exist as well. For example, driving on the left side of the road is a convention in the United Kingdom; woe betide those who ignore this convention and drive on the right side instead. Another convention governs commercial transactions: in some such transactions, the buyer pays before the seller delivers, while in others, the seller delivers first, and then the buyer pays the bill. Again, problems can arise if the two parties involved do not conform to the same convention.

Convention, linguistic or otherwise, is not the only coordination device used for the successful achievement of joint actions. However, the different coordination devices described by Lewis and Clark are quite different in their applicability. The device of *explicit agreement* is when a group explicitly agrees to employ a particular coordination device, such as the explicit definitions of 'joint action', 'convention' and 'arbitrary' given above. However, explicit agreement can only be used when there is a rich enough communication system in which to cast the explicit agreement. *Conventions* only emerge after regular use for a recurrent coordination problem. They are not conventions on the first, or even the second or third occasions of use; they must come to be adopted as such by a community and be part of the community's common ground. *Precedent* is also a non-conventional coordination device: it worked once, so I'll try it again, and hopefully it will work this time as well. But precedent also cannot be a coordination device the first time; it requires a precedent, of course. The coordination device that is most basic and most general is *joint salience*, which includes joint attention, described above. The individuals engaged in a joint action exploit properties of the situation whose salience is common ground by virtue of our human perceptual and cognitive faculties (e.g. a large moving object in the visual field), or by virtue of a common cultural heritage (e.g. in playing the cello sonata, the culturally defined preference for slowing the tempo at a major cadence in a piece of Western classical music, not to mention the culturally defined notion of a major musical cadence in the same tradition).

Language use exploits all of these coordination devices in addition to convention (Clark 1996: 77–81; Croft 2000: 99–104). All of these coordination devices, conventional and non-conventional, are general social cognitive abilities. Likewise, language is a joint action; but the ability to carry out joint actions is also a general social cognitive ability. Cognitive linguistics has not included these social cognitive abilities in its model of language. Cognitive linguistics must do so; but cognitive linguistics can maintain its principle that these social cognitive abilities are general, and not properties of a specific, innate language module, not even a socially-designed language module.

Finally, by embedding language in its social cognitive context, we can offer a more precise definition of the function of language, a necessary prerequisite for any functionalist model of language. Language is a (largely) conventional coordination device to solve the coordination problem of communication, which in turn is a coordination device to solve the coordination problem of successfully achieving any joint action that human beings wish to engage in (see Clark 1999). That joint action can be anything from the sharing of information to the overthrow of the government.

3.2 From symbols to a semiotic triangle

The second basic principle of cognitive linguistics outlined in §2 is that *grammar is symbolic, and thus meaning is an essential part of grammar*. This principle is actually implicit in Lewis's and Clark's definition of convention, including linguistic convention. A convention is a regularity in behavior, partly arbitrary, used as a coordination device to solve a recurrent coordination problem. The regularity in behavior is the signifier and the recurrent coordination problem is the signified, in symbolic or semiotic terms. Even Saussure's *l'arbitraire du signe* is replicated in Lewis's definition of convention as (partly) arbitrary. A convention is thus a pairing of form and meaning, just as in the cognitive linguistic principle.

But the cognitive linguistic principle leaves out a crucial part of the definition of convention, namely that the form-meaning pairing (the convention) is common ground in a community. That is, the symbol is shared among certain individuals. Clark recognizes this by redefining symbols as communal lexicons, so that *grammar consists of a 'semiotic triangle' of the form, the meaning, and the community in which the meaning is conventional*. This is the revised second basic principle. (I have chosen the term 'semiotic triangle' here, and the three elements of the semiotic triangle here are not the same three referred to by Ogden and Richards as the triangle of reference [symbol, thought, referent; Ogden and Richards 1947: 11] or Peirce's trichotomy of the sign [representamen or sign in speaker's mind, interpretant or sign in hearer's mind, and object; Peirce 1932: 135].)

An example of the necessity of including the third leg of the semiotic triangle in order to understand the nature of a linguistic symbol is the word form *subject*. In the community of linguists, a *subject* is the most prominent grammatical relation in the clause. In fact, further subcommunities of linguists have different meanings regarding exactly how the subject is the 'most prominent' grammatical relation, and regarding exactly what is a 'grammatical relation'. When we move to the community of psychologists, we find yet another meaning for *subject*, namely a person undergoing an experiment. In yet another

community, that of university students, at least British university students, *subject* means the area of study. Finally, to a community which we may define as the community of laypersons, *subject* means the topic of a conversation or a text (as in, *The subject of this book is the Vietnam war*). These examples can be multiplied indefinitely, making it clear that even in the case of one 'language', English, one must consider it to be not just a set of conventional symbolic units, but rather multiple sets of symbolic units that can be represented by semiotic triangles or communal lexicons (Clark 1998) (or as cognitive linguistics would prefer, communal constructicons).

If so, then how is English a 'language', as a set of shared linguistic conventions? Our understanding of 'language' must accommodate a fundamental fact about social structure, succinctly described by Dwight Bolinger in the following passage:

There is no limit to the ways in which human beings league themselves together for self-identification, security, gain, amusement, worship, or any of the other purposes that are held in common; consequently there is no limit to the number and variety of speech communities that are to be found in society.

(Bolinger 1975: 333, quoted in Wardhaugh 1992: 126)

Every society, here taken to be a group of individuals, consists of multiple communities, including multiple speech communities. For example, English society consists of communities including the communities of linguists, of psychologists, of university students, and so on. Moreover, all individuals are members of multiple communities. I am a linguist, and also a 'layperson', and a member of many other communities in which I participate. (One consequence of this is that a speech community cannot be defined in terms of a set of individuals; we turn to this problem in the next section.) Hence, all individuals have a repertoire of codes for use in the different communities to which they belong. I can write like a linguist, and talk like a layperson. I can also if necessary write somewhat like a psychologist to the extent that I am familiar with the practices of that community and the code they use (and by virtue of this ability, I am a peripheral member of that community).

These are all social facts which must be integrated into a symbolic approach to the nature of a language. The result of integrating this social, in fact fairly standard sociolinguistic, approach is that our definition of a language must be considerably modified. A language is now a heterogeneous entity, with its structure provided by social structure as well as by the meanings communicated. This is the contribution of the third leg of the semiotic triangle. Again, it is not incompatible with the basic principle of cognitive linguistics that we started with. Sociolinguists describe different ways of saying the same thing (Weinreich, Labov and Herzog 1968: 162). 'Ways of saying the same thing' are the same as symbolic units enriched by the specification of the community (social group) with which each way of saying something is associated; in other words, the semiotic triangle.

3.3 From encyclopedic meaning to shared meaning

The third basic principle of cognitive linguistics is that *meaning is encyclopedic*; that is, a speaker's representation of the meaning of a word or construction is the totality of knowledge and experience of the situations for which that word/ construction has been used.

Again, however, cognitive linguistics has taken a solipsistic perspective. When I use a word, I cannot assume that you can read my mind and possess my encyclopedic knowledge that guides my use of the word on this occasion. However, I can assume that much of my knowledge, including much rich detail, is common ground between us, and on the basis of that common ground, communication is largely successful. In other words, in order to understand how linguistic meaning works in communication, and hence how speakers represent linguistic meaning, we must alter the cognitive linguistic principle to the principle that *meaning is shared*. That is, our understanding of an utterance relies on shared knowledge, beliefs and attitudes about our world, both natural and cultural.

This is the notion of *common ground*, which we have already seen plays a central role in the achievement of joint actions and in the use of convention as a coordination device for joint actions including communication. Common ground is sometimes described as 'shared knowledge', but in fact it includes shared beliefs and even attitudes among people. For example, to understand and appreciate the bumper sticker *Friends don't let friends vote Republican* requires shared beliefs and attitudes as well as shared knowledge about Republicans. Clark and Carlson (1981/1992) further argues that common ground is exactly the notion of 'context' that is frequently invoked in analyzing the interpretation of words and constructions. When, as a linguist, I say something like 'When using a definite noun phrase in a sentence like *Bring me the chair*, the context tells the addressee which chair the speaker is referring to', I mean that shared knowledge between speaker and addressee about the room, the furniture in the room, the goals of the speaker, what has already been said in the conversation, and so on, allows the addressee to correctly identify which chair the speaker is referring to.

Of course, common ground is not easy to establish because of that fundamental fact that we can't read each other's minds. But we are both human beings with similar behaviors who live in the same world. This equally fundamental fact, which we as human beings have the ability to recognize (this is Tomasello's hypothesis that we have the ability to recognize others as intentional agents like ourselves), provides bases for common ground. Clark argues that common ground is founded on a shared basis, which has the following properties: it provides information to the persons involved that it holds; it indicates to each person that it provides information to every person that it holds; and it indicates the proposition in the common ground (Clark 1996: 94; for justification of this analysis, see Clark 1996: 94–100 and Clark and Marshall 1981/1992).

Note that all I possess as an individual is information justifying that a shared basis exists, from which I infer the existence of the common ground indicated by that shared basis (Clark 1996: 96). This is not infallible of course: we make mistakes about what is common ground from what we think is the shared basis, and some shared bases provide less justification for our belief about common ground than others. This point will be of direct relevance to my application of a social perspective to a cognitive linguistic issue in §4. But it is strong enough for joint action and coordination to succeed enough of the time, and thereby allow human society to exist.

Clark argues that there are two types of common ground that can be defined in terms of the bases that support them. Clark's taxonomy can be related to recent work in social theory that has been brought into sociolinguistic research.

The first type of common ground is what Clark calls personal common ground. Personal common ground is shared directly in face-to-face interaction by the persons. Personal common ground has two bases. The first is what Clark calls the perceptual basis: we share knowledge of what is in our shared perceptual field, as in the example of *Bring me the chair* mentioned above. The perceptual basis is established by joint attention (Croft 2000: 94). Something I perceive is not part of personal common ground if I don't have reason to believe that you are attending to it also. The second basis is what Clark calls the actional basis, but the term discourse basis would be more accurate (Croft 2000: 94). This is common ground shared by virtue of shared conversations. There are situations we have not experienced perceptually together, but when I report situations I have experienced to you in conversation, and vice versa, these become part of our personal common ground. Although we did not experience them perceptually together, we did experience the reporting of them linguistically together. The discourse basis thus involves joint attention (to the linguistic signal), as well as the common ground of a shared language (Croft 2000: 94).

Personal common ground, as the result of face-to-face interaction, can be related to the notion of a social network (Granovetter 1973, 1982), which is defined by direct interactions of individuals to varying degrees (network density) and in varying numbers of communities in which individuals interact (network multiplexity). The structure of social networks has been argued by Leslie Milroy to be instrumental for language maintenance and change (Milroy 1987; see also the discussion in Croft 2000: 169).

The second type of common ground is what Clark calls communal common ground. Communal common ground is shared by virtue of common community membership. For example, if I meet someone at a cognitive linguistics conference, and I can infer from the situation that he or she is a cognitive linguist, then I can infer shared knowledge, beliefs and attitudes simply by virtue of the fact that we are both cognitive linguists, even though we have never met face-to-face before and therefore have no prior personal common ground. Hence communal common ground is a powerful means for establishing shared meaning. Some communities are quite specialized, such as the community of cognitive linguists. But other communities are very broad and even all-encompassing. I can establish common ground with strangers by virtue of our both belonging to the community of Western urban industrialized society, the community of men, or ultimately the community of human beings on Earth, the community to which we all belong.

Clark proposes that the basis of communal common ground is shared expertise. I have learned how to be a cognitive linguist, and to do cognitive linguistics; that is my expertise as a member of this community. The person I have never met before at a cognitive linguistics conference has acquired the same expertise, and on that basis we have our communal common ground. I have also learned how to be a human being and about general human behavior; that is my expertise as a member of the all-encompassing community of human beings.

However, more recent work in sociolinguistic theory (e.g. Eckert 2000) has proposed that a better definition of community is in terms of communities of practice (Wenger 1998). I would like to suggest that shared expertise emerges from shared practice – one of Wenger's central insights. However, Wenger's definition of communities of practice is very different from Clark's definition of communities.

Wenger's goals and Clark's goals in defining communities are quite different. Clark aims for an all-encompassing approach to communities, while Wenger isolates a particular kind of community that he considers essential for the learning process. One of Clark's goals in defining a community is understanding how strangers can successfully communicate by virtue of communal common ground. Wenger's central concern, on the other hand, is how a group of individuals together negotiate meaning in the world via shared practice – 'meaning' in the sense of 'our ability to experience our life and the world as meaningful' (Wenger 1998:5), which I will call 'meaningfulness' to distinguish it from meaning in the linguistic semantic sense.

Wenger defines communities of practice in terms of properties that closely resemble Bratman's definition of joint action (shared cooperative activity) and the Lewis and Clark definition of convention. There is a mutual engagement among the individuals in the community of practice (Wenger 1998:73–77). These are clearly joint actions as we have described them above. Wenger states that a community of practice is ultimately defined by mutual engagement, that is, participation in joint actions (*ibid.*:73–74). Wenger discounts the social network model, defined according to him as the network 'through which information flows' (*ibid.*:74). But if communication is a coordination device for joint actions, and 'information flow' is the joint action of sharing information, then the social network model is compatible with Wenger's community of practice, as the social network is ultimately defined by joint actions. It is just that Wenger focuses, correctly in my view, on the joint actions rather than communication; the latter is simply the coordination device for the former.

The community of practice is further defined by a joint enterprise, negotiated by the members of the community (Wenger 1998:77–80). I would describe the joint enterprise as a joint purpose for the joint action, or perhaps the higher levels of the joint action: recall that joint actions can be nested in other joint actions. (These are the meshing subplans of the participants.) The joint enterprise can also be broadly construed as what Bolinger describes in the quote above as any purpose that leagues persons together, though it appears that Wenger intends a more lasting or structured purpose for his community of practice (*ibid.*: 125–126). Wenger emphasizes that the joint enterprise also entails mutual accountability (*ibid.*:81–82). This property is essentially Bratman's criterion of a commitment of mutual support, which leads to mutual accountability for that support.

Finally, the community of practice is defined by a shared repertoire for carrying out the joint enterprise (Wenger 1998:82–84). Wenger describes the shared repertoire of a community of practice as including 'routines, words, tools, ways of doing things, stories, gestures, symbols, genres, actions or concepts ... which have become part of its practice' (*ibid.*:83). This shared repertoire can be described as including the conventions employed by that community to solve the coordination problems in their mutual engagement (joint actions). The shared repertoire also includes the meshing subplans conventionally used to achieve the joint enterprise (which may be conventions in the Lewis and Clark sense only by unduly stretching the notion of coordination). Finally, the shared repertoire also appears to include what Clark would call the shared expertise, and may even constitute that shared expertise.

For Wenger, communities of practice are defined essentially by mutual engagement. Such communities develop personal common ground since they possess shared perceptual and discourse bases. But they also develop shared expertise through the practice that defines them. Wenger is certainly correct that shared expertise emerges from – is learned through – shared practice. Yet Clark extends the notion of shared expertise to something that can exist between strangers, even as fellow humans, where mutual engagement is absent. Is this difference between Wenger's and Clark's views of communities irreconcilable?

In Wenger's definition of communities of practice, the repertoire for carrying out the joint enterprise can be shared by members of a community that have never met; but that is all that is shared among them (Wenger 1998: 291, footnote 1 to Chapter 5). This is not enough to define a community of practice for Wenger. Wenger argues (*ibid.*: Chapter 3), that 'communities of practice can be thought of as shared histories of learning'. I agree that learning – the acquisition of shared expertise – can only take place through mutual engagement in communities of practice. Nevertheless, there are two aspects of Wenger's definition which can be stretched and lead to communities defined by shared expertise that include strangers, as Clark intends. (And these larger communities are more structured than Wenger's overlapping and intersecting communities of practice which he calls 'constellations' [Wenger 1998: 126–131].)

First, the notion of joint enterprise can be construed broadly. For example, two cognitive linguists partake in the joint enterprise of explaining linguistic phenomena in accordance with the basic principles presented in §2, even if they are working in countries on opposite sides of the globe. More specific joint enterprises, such as analyzing the perceptual metaphors of Serbian from a cognitive linguistic perspective, are not necessarily shared by all members of the cognitive linguistics community. Nevertheless, they may be shared by members of the more exclusive Serbian metaphor cognitive linguistics community who could be strangers. Wenger would not accept this definition of 'joint' because mutual engagement is absent: you and I are laboring separately in analyzing the perceptual metaphors of Serbian.

Any member of any community in Clark's sense gains their communal expertise by belonging to a community of practice in Wenger's sense. For example, I studied and learned from other cognitive linguists, and collaborated with other cognitive linguists, and mutually engage with cognitive linguists by talking with them, giving lectures, writing publications, and by listening to other cognitive linguists and reading their publications. So the expertise that I share with cognitive linguists who I don't know was acquired by being part of a community of practice with cognitive linguists who I do mutually engage with, either in face-to-face interaction or other media (email discussion, written publications).

Wenger describes communities of practice as shared histories of learning, as quoted above. But sharing can be indirect as well as direct. I engage mutually with Len Talmy, for example, through conversation and his publications, and gain a certain degree of shared expertise which is part of the shared expertise of cognitive linguistics. I then engage mutually with my semantics students, who may never meet or even read Talmy. My students and I then acquire some shared expertise, which is indirectly shared with Talmy. This indi-

rectly shared expertise between my semantics students and Talmy would form a basis for common ground should my students ever meet Talmy or read his work. Likewise, I engage mutually with Len Talmy as described, and so do you, but not in contact with me. When we meet and correspond as cognitive linguists, we can form a basis for common ground by virtue of our shared expertise acquired by engagement with the cognitive linguistics canon, including Talmy's contributions or Talmy himself.

The crucial insight is that the acquisition of shared expertise is a shared historical process. While Wenger would restrict the community of practice through directly shared historical expertise, the historical nature of the process and the continually changing but overlapping membership of communities, and their spread and divergence through the world, means that expertise can be shared indirectly, from A to B and B to C (where A and C do not interact), or from A to B and A to C (where B and C do not interact). It is difficult if not impossible to draw a sharp line between direct and indirect sharing. Wenger's case study of a community of practice is a group of medical claims processors in an insurance company office. In this community, turnover is high and yet there is continuity as A departs and later C arrives and is trained by B who worked with A. Of course, indirectly shared history provides a less strong shared basis than directly shared history; but the difference is a matter of degree, not kind. The evolutionary biologists would call this history of shared practices, direct and indirect, a lineage; a social scientist would call it a cultural tradition. I will argue later that this evolutionary perspective is crucial for understanding the nature of language.

We can preserve Wenger's narrow notion of a community of practice as the locus of learning, which I interpret as the acquisition of shared expertise. But even when direct mutual engagement is absent, the historical nature of the learning process and hence the formation of the community of practice means that expertise can be shared beyond mutual engagement. There is a sense in which individuals not directly mutually engaged still have a joint enterprise (not Wenger's sense, I should caution). There is a history of joint actions – mutual engagements – that make up the larger-scale joint enterprise. This history forms lineages and spawns branching lineages. These lineages form larger groupings that share expertise and represent communities in Clark's sense. But they are ultimately rooted in mutual engagement in local communities of practice in Wenger's sense: this is how we join a community (in Clark's sense) and inherit its shared expertise.

Our directly shared experiences, and the shared expertise that emerges from the cultural tradition of a community of practice, are encyclopedic in the cognitive linguistic sense. Nothing in common ground is excluded a priori from contributing to linguistic meaning or its representation in a speaker's grammar. But what matters at least as much is that the meaning is shared, and who it is shared with, since language's function is to serve joint actions.

3.4 From meaning as construal to construal for communication

The fourth and final basic principle of cognitive linguistics is that *meaning involves conceptualization (construal)*. A conceptualization process, or construal operation, has three characteristics. First, there are alternative possible construals of a particular scene. For

example, I can describe the same scene with the alternatives of a mass noun *foliage* or a bare plural count noun *leaves*; these represent two different construals of the phenomenon (alternative structural schematizations, in Talmy's terms: Talmy 2000, Chapter 1). Second, one must take a particular construal of a scene; one cannot avoid construing the scene in a particular way. We must choose the homogeneous construal of the mass noun *foliage*, or the aggregate construal of the bare plural count noun *leaves*, or possibly some other construal offered by a more complex linguistic expression. But we cannot avoid construing the scene in one way or another. Finally, no construal has an a priori privileged status as the "best" or the "correct" construal of a scene. It is not "better" or "more correct" to construe the scene as homogeneous (*foliage*) or an aggregate (*leaves*). Either is valid a priori, although one may be considered more useful on a particular occasion of use, depending on the goals of the interlocutors.

Numerous construal operations have been identified by many cognitive linguists including Talmy, Lakoff and Langacker. The construal operations can be grouped together and linked to cognitive psychological processes. In Croft and Cruse (2004, Chapter 3) we identify four general psychological categories of construal operations: attention, comparison, perspective and gestalt. Many cognitive linguists, in particular Talmy and Langacker, argue that various grammatical elements or constructions have as their chief 'meaning' a particular conceptualization of the experience denoted by the lexical item found with the grammatical element or construction.

There is another functionalist linguistic tradition that instead analyzes the meaning of various grammatical constructions in terms of their 'discourse function', 'information structure' or 'information packaging'. These two traditions, the cognitive linguistic and the discourse functional, are sometimes perceived as representing opposing approaches to the analysis of grammar. But they can be reconciled. From the cognitive linguistic perspective, this requires the recognition that construal is not simply a cognitive process. Construal is always for a purpose in a communicative act. That is, *meaning involves construal for the purpose of communication*. This is the social cognitive reformulation of the final basic principle of cognitive linguistics.

Exactly how this is the case requires a brief analysis of the linguistic act. Here we can turn to speech act theory, as originated by Austin (1962) and developed by Searle (1969) and Clark (1992, 1996). The aspect of speech act theory that is most lasting and relevant to us is not the classification of illocutionary acts but the model of what actually goes on when there is a successful joint action involving speech. Again, Clark has provided a reinterpretation of the Austin-Searle model of speech acts as joint actions, as well as contributing an important missing link.

Clark describes a speech act as simultaneously involving four levels (Clark 1996: 148–153), three of which were recognized by Austin and Searle. The first level is the utterance act (Searle 1969: 23–24; compare Austin's phonetic act [Austin 1962: 92, 95]): the speaker produces an acoustic signal, and the addressee attends to it. That is, joint attention is established on the acoustic signal created by the speaker. This is a joint action in that both speaker and addressee must perform their parts and their parts must be coordinated: the speaker produces a clear and audible signal, and the addressee pays attention to the signal.

If either individual fails in their action, then the utterance act will fail. This is the level recognized by linguists as the phonological structure of a language.

The second level is the propositional act (Searle 1969: 23–24; compare Austin's phatic and rhetic acts, which separate form and meaning/reference [Austin 1962: 92, 95–100]): the speaker formulates his utterance in a particular language with its particular constructions, and the addressee recognizes the formulation, both the language it is formulated in and the words and constructions from that language which the speaker has employed. The propositional act level is the locus of linguistic convention as described in §3.1; that is, it is the level where meaning in the sense of the signified in a linguistic sign or symbol is found. This level is also the locus of construal, since construal is a property of the meanings of words and constructions in a language. This is also the level recognized by linguists as the grammatical (morphosyntactic) structure of a language.

The propositional act is of course simultaneous with the utterance act. In fact, they are the same action in one sense – we are talking about just one thing that is happening at that moment – but not the same action in another sense. For example, the utterance act can occur in a different medium (writing). The two acts are linked to each other in an asymmetric fashion. The propositional act cannot succeed unless the utterance act succeeds: if I don't hear you clearly, I won't recognize what you are saying. However, the utterance act could succeed while the propositional act fails: I could hear you clearly, but not recognize the formulation because it is in a language unknown to me. Clark describes this asymmetrical relationship between the 'two acts in one' as an 'action ladder' (Clark 1996: 147–148). Successful achievement of the 'upper' propositional act presupposes successful completion of the 'lower' utterance act (upward completion); it therefore also provides evidence that the lower act is complete (downward evidence).

The third level in the speech act action ladder is the informative act (Clark and Carlson 1982/1992): the speaker is presenting the intention she wishes the addressee to recognize, and the addressee understands the speaker's intention. This level had been overlooked before Clark and Carlson, but it is crucial because of the fundamental fact of social life that we cannot read each other's minds. The informative act is what the speaker is normally trying to accomplish with her formulation in the propositional act: making public her intention. Again, there is an asymmetric relationship: the informative act can succeed only if the propositional act is successful. This is the level described as 'communication', and also the level at which the Gricean concept of meaning is defined, roughly that the speaker means something by intending that the hearer should recognize her intention (Grice 1948/1989; see Clark and Marshall 1981/1992: 19 and Clark 1996: 129–130 for a revision of Grice's definition). The asymmetric relationship between propositional act and informative act is precisely the one described between linguistic convention and communication in §3.1: linguistic convention functions as a coordination device for the joint action of communication.

The fourth and final level is the joint action that the speaker is attempting to engage the addressee in: the speaker proposes the 'joint project', as Clark calls it, and the addressee takes it up in one way or another. This is the level of the illocutionary act as it is normally described in speech act theory. But speech act theory did not recognize its joint character (at least at first). Also, speech act theory attempted to construct an exhaustive taxonomy

of illocutionary acts. This proved to be impossible because this level crosses into social interaction in general: a linguistic communicative act can function as a coordination device for the achievement of any joint action that members of the society wish to engage in. Again, the speaker and addressee cannot engage in a joint action unless the informative act was successful; that is, the speaker's intention has been successfully recognized by the addressee. This is the respect in which communication (the informative act) is serving as a coordination device for the successful execution of a joint action.

This model of the 'four acts in one' that takes place in any linguistic intercourse identifies the locus of construal at the propositional act level. But it also demonstrates that this is only one level in the speech act, and serves for the achievement of the higher levels in the action ladder (the informative act and the joint project/action). The identification of construal as an essential part of symbolic meaning, and relating it to cognitive psychological processes, is a major contribution of cognitive semantics. But it must be situated in the larger model of language, communication and joint action in order to understand why it exists and how it is used by speaker and listener.

3.5 Summary

In this section, I have argued that to create a social cognitive linguistics, we must replace or modify the four basic principles of cognitive linguistics in §2 with the following four basic principles:

Grammatical structures and processes in the mind are instances of general social cognitive abilities as well as individual cognitive abilities

Grammar consists of a semiotic triangle of the form, the meaning, and the community in which the meaning is conventional

Meaning is shared as well as encyclopedic

Meaning involves construal for the purpose of communication

These slogans are intended to evoke a coherent model of the role of language in social interaction and the role of social interaction in language developed by linguists, philosophers, psychologists and social theorists. Language functions as a conventional coordination device for communication, which in turn is a coordination device for joint actions, which can be thought of as the glue holding society together. These form three of the four levels of the action ladder of speech acts. (The fourth, 'lowest' level is the physical realization of linguistic conventions in utterances, typically as sound structure but also in gestural sign language or in writing, which provide the perceptual basis for successful communication.) Convention is one coordination device among others, of which the most important is joint salience, based on the social cognitive abilities of joint attention and recognition of others as intentional agents. Joint attention and linguistic conversations provide the basis for personal common ground, as well as making joint actions possible. Moreover, cultural traditions of joint actions (mutual engagement) for a joint purpose (enterprise) define communities, and provide the basis for communal common ground. Communities over-

lap within societies, and all of us are members of multiple communities, which means that our language is multiplex, reflecting the structure of the society in which it is spoken.

4. The interplay of cognition and social interaction: A case study

In the last part of this chapter, I describe how moving towards a social cognitive linguistics offers new perspectives into traditional cognitive linguistic issues, as well as bringing new issues into the purview of (social) cognitive linguistics. In particular, I will bring linguistic variation into a social cognitive perspective, and reexamine the phenomenon of construal in this light.

The phenomenon of construal has been used to account for the meanings of grammatical forms and constructions, so that e.g. *foliage* and *leaves* impose different construals on the experience by virtue of occurring in the bare singular (mass noun) and bare plural constructions respectively. In this respect, then, *foliage* and *leaves* are not synonymous: among other things, they represent different conceptualizations of an experience.

I argued in §3.4 that construal is always for the purpose of communication in a particular occasion of use. If this is true, then in order to understand the nature of construal, we must consider construal in the light of the use of language for communication in joint actions, the nature of the communities to which the interlocutors belong, and the nature of common ground and the evidence each of us has for it. When we look at construal in this way, a somewhat different picture of it emerges than when we think of construal as an internal mental process of conceptualization.

In order to investigate the role of construal in communication, one must examine language in use, that is, the verbalization of experience (Chafe 1977a, b). In particular, we can investigate variation in multiple verbalizations of a particular experience in order to determine what role, if any, construal is playing in alternative verbalizations of experience. This is most easily done in an experimental paradigm. A rich source of data is the Pear Stories project (Chafe 1980). A six-minute film without any spoken language, designed by Chafe and colleagues, was shown to Berkeley undergraduates, who then described the film to a third party. This is a reasonably controlled study: all of the speakers have a similar sociolinguistic background (similar age and educational level), they all saw the same sequence of scenes, and they all verbalized what they saw in a similar, albeit artificial, environment. (The Pear film was also shown to speakers of other languages.)

The papers published in Chafe (1980) mostly focus on sociolinguistic and discourse functional aspects of the verbalization process. I used the corpus of twenty English narratives to examine the semantics of verbalization; that is, how the scenes and their elements were verbalized in words and constructions (Croft to appear). Here I will discuss one example from my study and its significance for issues in a social cognitive linguistics.

The following example (Table 1) represents the verbalizations of part of one scene in the film (scene A5 in Croft to appear; the number x,y at the left refers to intonation unit y of speaker's x's narrative).

Table 1. The verbalization of scene A5 in the Pear Stories film

1,11	and he'd drop them into his [.25] thing,
3,7	[3.15 um [2.35] um [.35]] picks the pears and puts them in a [.45] in um [.4] these baskets that he has.
5,10	.. filled with [.8] and he's filling them with pears.
6,8	and putting them in a . . white apron,
8,12	[.6 [.1] A--nd] he [1.0] fills his-- . . thing with pears,
10,27	and then he'll . . he'll stuff them in that,
11,9	[.95] puts them in . . his apron,
12,21	and he put them in an apron /that he had/,
13,9	and putting it in his . . apron,
15,11	[.15] but he was also he also had an apron on and he was filling /those/ up.
17,20	[.45] putting them in his apron,
18,14	and putting them in [.7 . . . {breath} . . .] his apron.

The data demonstrate that even with one part of one scene, each speaker verbalizes that part slightly differently. If we compare other similar scenes, “putting” scenes, in the Pear Stories data, we find that each of them is variable in its verbalization across speakers. This fact must be confronted by any theory of language.

The standard cognitive linguistic analysis of this variation in verbalization is in terms of construal. Each speaker construes scene A5 in the Pear Stories film slightly differently, and the verbalizations thus represent different construals of the scene. In other words, for each speaker, the scene actually being verbalized is slightly different, because what is being verbalized is a construal of the scene. That is, the different verbalizations given above are not really verbalizations of the same scene. In this analysis, variation is explained away. Each speaker-conceptualized scene is different and precisely verbalized. There is thus no variation in verbalization. This analysis conforms to the widespread belief that there is no real synonymy in language: every grammatical difference corresponds to a semantic difference, often a difference in conceptualization.

However, there are reasons to believe that the standard cognitive linguistic analysis is incorrect, and these reasons are found in the social interactional nature of language. The verbalizations illustrated above were all communicative joint actions, and therefore we must consider not only the speaker but the hearer of the verbalization.

The speaker who wishes to verbalize a scene like A5 has a range of alternative conceptualizations of the scene available to her, according to the various conceptualization processes described by cognitive linguists. In fact, the range of alternative conceptualizations is enormous, since various facets of attention, comparison, different kinds of perspectives and alternative gestalts of the scene and its parts are all available. However, the hearer also has the same huge range of alternative conceptualizations available to him. Now the fundamental fact of communication as a joint action enters: the hearer cannot read the speaker's mind, and therefore cannot be certain of the precise construal intended by the speaker (§3.1).

How does the hearer know which of the huge number of construals of the scene is the one intended by the speaker? The speaker chooses one verbalization of the scene out of

the huge number of possible verbalizations, each verbalization supposedly corresponding to a possible construal of the scene. The speaker's choice of words and constructions to verbalize is based on her prior exposure to and use of those words and constructions in other communicative acts. On this basis, the speaker has schematized a construal of real world situations that can be applied to the scene in the current communicative event. But the same applies to the hearer. His knowledge of the words and constructions that he hears produced by the speaker is based on his own past exposure to and use of them in other communicative acts. Yet the hearer's past history of use of the words and constructions is different from that of the speaker. Thus, the hearer's schematic conceptualization will differ from the speaker's. Thus even when hearing familiar words and constructions, the hearer cannot be certain that the conceptualization he thinks they represent is the conceptualization that the speaker thinks she is conveying.

It doesn't stop there. No two experiences are identical. Each speaker's experience is in some ways new and different from all previous experiences, at the same time that it is also similar. Therefore any choice of words and constructions will not precisely characterize the construal of the scene being communicated anyway. Every verbalization relates a current experience (construal of a scene) to prior experiences. This is the notion of recurrence in the recurrent coordination problem that is solved by the convention of using a linguistic form (§3.1). But the current situation, or more precisely the part of the current situation verbalized by a particular word or construction, is unique and hence not quite the same coordination problem. Therefore, using a particular choice of word/construction in the current situation construes it as being like certain prior situations in the language user's experience, but also modifies the linguistic system by sanctioning the application of the linguistic forms to the new experience.

Thus, there is a fundamental indeterminacy in the construal of a scene by a speaker and its interpretation by the hearer in a communicative act. Moreover, this is true of every communicative act: not just the present communicative act, but every prior communicative act which forms the basis of the speaker's and hearer's understanding of the meaning of the words and constructions used in the current utterance. Thus, alternative construals provided by alternative verbalizations cannot be precise.

Instead, both speaker and hearer must fall back on their experience of the current scene, as well as their prior experiences – in other words, their common ground (§3.3). If both interlocutors witnessed the scene, they can use the perceptual basis (joint attention and salience). In the case of recalled experience, as with the Pear Stories (the speaker was told that her interlocutor had not seen the film), speaker and hearer must rely on shared expertise of the meanings for which the linguistic forms in the utterance have been used. As I noted in §3.3, what each person has is information that a shared basis for the common ground holds. Linguistic conventions are part of the interlocutor's communal common ground; their use provides a shared basis for understanding of the scene. But the information they contain is of varying quality (e.g. how familiar I am with the words and with how they have been used). And the linguistic conventions, though part of a common cultural tradition, are likely to be only indirectly shared expertise, unless the interlocutors are part of the same community of practice in Wenger's sense (see §3.3).

Table 2. Alternative verbalizations of scenes with *put*

	<i>Put</i>	Other verb	Other verbs used
A5	8	2	<i>drop, stuff</i>
A7	9	13	<i>empty, dump, tumble, drop, place, deposit</i>
C5	15	-	
E5	8	4	<i>load, throw, toss, pour</i>
G3	1	6	<i>deposit, dump, empty, unload</i>

The argument in the preceding paragraphs indicates that one cannot put too great a precision on the shared semantics of linguistic forms. This is in fact not a bad thing, because the degree of variation in verbalization of the same scene turns out to be quite great. And above all, languages change, and the possibility of change is ultimately due to the indeterminacy of verbalization.

Consider the choice of a verb to describe the event in the verbalization of scene A5 and related scenes in the Pear Stories. Of course, choice of verb is only one part of the verbalization of the scene, so focusing on the verb abstracts away from much of the variation in verbalization. It is the verb choice combined with other linguistic choices that leads to the uniqueness of every verbalization as noted above. Table 2 indicates the alternative verbalizations to *put* in all of the scenes in the Pear Stories where at least one speaker used *put*.

As can be seen, in four of the five scenes, there was variation in the verb chosen. In scene C5, only *put* is used. Unlike the other scenes, the putting subevent is part of a sequence of events that required verbalization by multiple predicates; the verbalization of the entire scene is of course highly variable. (It may also be significant that the other scenes involve putting pears or a pear into a basket, whereas scene C5 involves putting the whole basket onto a bicycle. At any rate, the exclusive use of *put* in scene C5 is anomalous in the context of all the data analyzed in Croft to appear.)

The verb *put* is sometimes called a light verb because it is used in such a broad range of situation types. The other verbs chosen by different speakers to verbalize the event are used for more specific situation types, typically expressing the manner of placement. The verbs are sometimes semantically incompatible. For example, an introspective semantic analysis of *drop* and *stuff* would not allow them to be used to describe the same scene, since dropping is assumed to entail letting go (Talmy's onset letting causation; Talmy 2000: 418) and stuffing involves continued application of force (Talmy's extended causation). Yet both *drop* and *stuff* are used for scene A5. This may be due to indeterminacy in recall; but given what I have just argued, a language user cannot distinguish a priori between indeterminacy in recall and indeterminacy in verbalization.

Thus, we observe that a range of verbs are used for putting events in the Pear Stories narratives, not just different putting events in different scenes of the film, but also for almost every individual scene in the film. This variation in the morphosyntactic verbalization of (construed) semantic scenes is essentially the same as the variation we see in the phonetic realization of phonemes. In phonology, Ohala (1989) and others after him have argued that 'sound change is drawn from a pool of synchronic variation'. Here, I argue that grammatical change is drawn from a pool of synchronic variation as well: the variation in

Table 3. Specific verbs in the Pear Stories as etymological sources of *put* (Buck 1949, §12.12)

Pear Stories	Indo-European light verb	Source/related verb in older language
<i>throw, toss</i>	Mod. Greek <i>vazo</i>	Anc. Greek <i>bállō</i> 'throw', occasionally 'put'
	French <i>mettre</i> , Italian <i>mettere</i>	Latin <i>mittere</i> 'let go, throw', Late Latin 'put'
	Mod. Irish <i>cuirim</i>	Old Irish <i>cuirim</i> 'throw, put'
<i>stuff</i>	English <i>put</i>	Old English <i>potian</i> 'thrust, push'
<i>place*</i>	Dutch <i>plaatsen</i>	Dutch <i>plaats</i> 'place (n.)'

* Already conventionalized with 'put' meaning in English

verbalization. In Croft (to appear), I tested this hypothesis for all the light verbs in the Pear Stories narratives as well as a host of grammatical function words and constructions. For the light verbs, I compared the more specific verbs used for scenes also expressed by light verbs such as *put* to the etymological sources of *put* verbs in Indo-European (Buck 1949). The result for *put* is given in Table 3.

The hypothesis is overwhelmingly confirmed in the data. The alternative, less grammaticalized verbalizations of scenes are also the etymological sources of the more grammaticalized forms used to verbalize the same scenes. This is possible only if language users interpret a range of words or constructions to be more or less alternative verbalizations of the same scene. The alternative verbalizations are interpreted as alternative ways of saying the same thing, and so the less grammaticalized form can then be propagated as the new convention of the speech community – that is, become grammaticalized.

5. Conclusion: An evolutionary, social cognitive model of language

In the preceding section, I have shown how variation is fundamental to grammar, even in the mapping from meaning to form, from experience to utterance. Variation is always liable to become change in progress. In other words, a social cognitive linguistics is also a dynamic, evolutionary linguistics.

Human beings engage in joint actions. This is a major part of what it means to be human, and is rendered possible by the general social cognitive abilities that humans have (compare Tomasello 2008). They succeed in their joint actions by employing the resources in the cultural traditions of their communities of practice. In order to coordinate these joint actions, human beings communicate. But since they cannot read each other's minds, they must coordinate their communicative acts by using language, which is a primary resource of their cultural tradition (and is part of their community of practice).

However, every situation is unique. Thus each linguistic communicative act represents a construal of the current joint enterprise as a recurrence of prior joint actions, via the replication of linguistic structures in the current utterance. This is an active process and is also a joint action: how the hearer interprets an utterance in relation to the situation being communicated is as important as how the speaker intended it. Fortunately, continued interaction between persons strengthens the joint interpretation of an utterance. If you say something to me, and you are not sure how I interpreted your utterance, then my utterance in response to you will tell you a lot about how I interpreted your utterance;

and you can negotiate the interpretation of your original utterance with me if you think it necessary. This is a familiar process to academics, where the interpretation of theoretical terms and concepts is difficult and intensely debated. It is also familiar to teachers and students in the interaction via lecture, discussion and student coursework and exams: it is through this process that teacher and student come to some degree of joint interpretation (or learning, as it is often called; and it is not always just learning on the student's part).

The fact that each language user has her own unique, if partly shared, history of language use, that each situation is unique, and the fact that we cannot read each other's minds, means that there is a fundamental indeterminacy of construal in conversation, although common ground means that communication is possible and often successful. One should also not lose sight of the fact that from the point of view of language's function as a coordination device for communication which is in turn a coordination device for joint action, absolute precision is not necessary for success in the function of language in the real world. As a consequence, there is a high degree of variation in verbalization of similar situations in a single speech community (and even by a single speaker), as we observed in the example in §4. We can describe this as *first order* variation (Croft 2006: 98–99).

The language users in a society will therefore have been exposed to multiple variant verbalizations for a particular meaning. The variants might come to be socially construed as linguistic resources for a particular community in the society, possibly a specific community of practice, and then valued or exploited as such. That is, the variants might come to be construed as variants of a sociolinguistic variable, with each variant now indexed for a particular community. The variants have now moved from being first order variants, the product of the indeterminacy of communication, to *second order* variants, indexes of community identity.

Then the dynamics of social structure can lead to the propagation of some variants and the extinction of others. The end result of this process will be language changes, which cause typological (crosslinguistic) variation in the case where one variant is propagated in one society and other variant is propagated in another society, both derived ultimately from a common ancestor society and language. This crosslinguistic variation is *third order* variation.

This is an evolutionary model of language (Croft 2000, 2006). It unifies the analysis of grammatical structure, function and variation at three different levels. The first is the basic social-interactive level where first-order variation is generated, and is the foundation of social cognitive linguistics. The second is the socially-indexed level. (Recall that symbols are really semiotic triangles of the form, meaning and the community indexed by the form-meaning pairing.) This is traditionally thought of as the realm of sociolinguistics, but it must be understood in the context of the first level of analysis of language use. The third is the level of crosslinguistic diversity, traditionally the area of typological research. Many if not most typologists seek functional-cognitive – perhaps I should say social cognitive – explanations for patterns of crosslinguistic variation. The integrated model outlined here shows how crosslinguistic diversity is a reflection of patterns ultimately rooted in social cognitive interactions between human beings (see also Croft 2001, 2003).

Language is a fundamentally heterogeneous, indeterminate, variable, dynamically unfolding phenomenon, just like the human society it constitutes a part of. Individual

cognition plays an essential though instrumental role in the dynamics of language. since human beings use their minds to interact with others in a way that appears unique to the human species. By extending cognitive linguistics to social cognition, it can play a role in advancing our understanding of the nature of language's role in humanity.

References

- Austin, J. L. 1962. *How to Do Things with Words*. Cambridge, MA: Harvard University Press.
- Bolinger, D. 1975. *Aspects of Language*. New York: Harcourt Brace Jovanovich.
- Bratman, M. 1992. Shared cooperative activity. *The Philosophical Review* 101: 327–341.
- Bratman, M. 1993. Shared intention. *Ethics* 104: 97–113.
- Bratman, M. 1997. I intend that we. In *Contemporary Action Theory, Volume II*, G. Holmström-Hintikka & R. Tuomela (eds), 49–63. Dordrecht: Kluwer.
- Buck, C. D. 1949. *A Dictionary of Selected Synonyms in the Principal Indo-European Languages*. Chicago: University of Chicago Press.
- Chafe, W. 1977a. Creativity in verbalization and its implications for the nature of stored knowledge. In *Discourse Production and Comprehension*, R. Freedle (ed.), 41–55. Norwood, NJ: Ablex.
- Chafe, W. 1977b. The recall and verbalization of past experience. In *Current Issues in Linguistic Theory*, P. Cole (ed.), 215–246. Bloomington: Indiana University Press.
- Chafe, W. 1980. *The Pear Stories*. New York: Ablex.
- Clark, H. H. 1992. *Arenas of Language Use*. Chicago and Stanford: University of Chicago Press and the Center for the Study of Language and Information.
- Clark, H. H. 1996. *Using Language*. Cambridge: Cambridge University Press.
- Clark, H. H. 1998. Communal lexicons. In *Context in Language Learning and Language Understanding*, K. Malmkjær & J. Williams (eds), 63–87. Cambridge: Cambridge University Press.
- Clark, H. H. 1999. On the origins of conversation. *Verbum* 21: 147–161.
- Clark, H. H. & Carlson, T. B. 1981. Context for comprehension. In *Attention and Performance IX*, J. Long & A. Baddeley (eds), 313–330. Hillsdale, NJ: Lawrence Erlbaum Associates. (Reprinted in Clark 1992, 60–77.)
- Clark, H. H. & Carlson, T. B. 1982. Hearers and speech acts. *Language* 58: 332–373. (Reprinted in Clark 1992: 205–247.)
- Clark, H. H. & Marshall, C. R. 1981. Reference and mutual knowledge. In *Elements of Discourse Understanding*, A. K. Joshi, B. L. Webber & I. A. Sag (eds), 10–63. Cambridge: Cambridge University Press. (Reprinted in Clark 1992, 9–59.)
- Croft, W. 2000. *Explaining Language Change: An Evolutionary Approach*. Harlow, Essex: Longman.
- Croft, W. 2001. *Radical Construction Grammar: Syntactic Theory in Typological Perspective*. Oxford: Oxford University Press.
- Croft, W. 2003. *Typology and Universals*. Cambridge: Cambridge University Press.
- Croft, W. 2006. The relevance of an evolutionary model to historical linguistics. In *Competing Models of Linguistic Change: Evolution and Beyond*, O. Nedergård Thomsen (ed.), 91–132. Amsterdam: John Benjamins.
- Croft, W. To appear. The origins of grammaticalization in the verbalization of experience. *Linguistics*.
- Croft, W. & Cruse D. A. 2004. *Cognitive Linguistics*. Cambridge: Cambridge University Press.
- Eckert, P. 2000. *Linguistic Variation as Social Practice: The Linguistic Construction of Identity in Belten High*. Oxford: Blackwell.
- Evans, V. & Green, M. 2006. *Cognitive Linguistics: An Introduction*. Mahwah, N.J.: Lawrence Erlbaum Associates.

- Fillmore, C. J. 1982. Frame semantics. In *Linguistics in the Morning Calm*, The Linguistic Society of Korea (ed.), 111–137. Seoul: Hanshin.
- Fillmore, C. J. 1985. Frames and the semantics of understanding. *Quaderni di Semantica* 6: 222–254.
- Fillmore, C. J., Kay, P. & O'Connor, M. C. 1988. Regularity and idiomaticity in grammatical constructions: The case of *let alone*. *Language* 64: 501–538.
- Goldberg, A. E. 1995. *Constructions: A Construction Grammar Approach to Argument Structure*. Chicago: University of Chicago Press.
- Goldberg, A. E. 2006. *Constructions at Work: The Nature of Generalization in Language*. Oxford: Oxford University Press.
- Granovetter, M. 1973. The strength of weak ties. *American Journal of Sociology* 78: 1360–1380.
- Granovetter, M. 1982. The strength of weak ties: A network theory revisited. In *Social Structure and Network Analysis*, P. V. Marsden & N. Lin (eds), 105–130. London: Sage.
- Grice, H. P. 1948/1989. Meaning. In *Studies in the Way of Words*, P. Grice (ed.), 213–223. Cambridge, MA: Harvard University Press.
- Langacker, R. W. 1987. *Foundations of Cognitive Grammar, Vol I: Theoretical Prerequisites*. Stanford: Stanford University Press.
- Lewis, D. 1969. *Convention*. Cambridge, MA: MIT Press.
- Milroy, L. 1987. *Language and Social Networks*. Oxford: Blackwell.
- Ogden, C. K. & Richards, I. A. 1947. *The Meaning of Meaning*. New York: Harcourt, Brace and Company.
- Ohala, J. 1989. Sound change is drawn from a pool of synchronic variation. In *Language Change: Contributions to the Study of its Causes*, L. E. Breivik & E. H. Jahr (eds), 173–198. Berlin: Mouton de Gruyter.
- Peirce, C. S. 1932. Ground, object and interpretant. *Collected Papers of Charles Sanders Peirce, Volume 2: Elements of Logic*, C. Hartshorne & P. Weiss (eds), 134–155. Cambridge, MA: Harvard University Press.
- Searle, J. R. 1969. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge: Cambridge University Press.
- Sinha, C. 1999. Grounding, mapping, and acts of meaning. In *Cognitive Linguistics: Foundations, Scope and Methodology*, T. Janssen & G. Redeker (eds), 223–255. Berlin: Mouton de Gruyter.
- Talmy, L. 2000. *Toward a Cognitive Semantics, Volume 1: Concept Structuring Systems*. Cambridge, MA: MIT Press.
- Tomasello, M. 1999. *The Cultural Origins of Human Cognition*. Cambridge, MA: Harvard University Press.
- Tomasello, M. 2003. *Constructing a Language: A Usage-Based Theory of Language Acquisition*. Cambridge, MA: Harvard University Press.
- Tomasello, M. 2008. *The Origins of Human Communication*. Cambridge, MA: MIT Press.
- Verhagen, A. 2005. *Constructions of Intersubjectivity: Discourse, Syntax, and Cognition*. Oxford: Oxford University Press.
- Wardhaugh, R. 1992. *An Introduction to Sociolinguistics*. Oxford: Blackwell.
- Weinreich, U., Labov, W. & Herzog, M. I. 1968. Empirical foundations for a theory of language change. In *Directions for Historical Linguistics*, W. P. Lehmann & Y. Malkiel (eds), 95–195. Austin: University of Texas Press.
- Wenger, É. 1998. *Communities of Practice: Learning, Meaning and Identity*. Cambridge: Cambridge University Press.

Cognitive and linguistic factors in evaluating text quality

Global versus local?*

Ruth A. Berman and Bracha Nir

1. The elusive notion of (expository) text quality

The study addresses two interrelated issues in the domain of (developing) text construction: the interplay between cognition and language in the integration of bottom-up and top-down cognitive processes, on the one hand, and the question of whether and how these relate to local linguistic expression as compared with global discourse organization, on the other. Our general goal is to throw light on the intuitively viable, yet elusive, notion of *text quality*. Intuitively, the notion makes good sense and is accessible to both non-expert readers and to expert evaluators, as demonstrated by the high inter-rater agreement observed in psycholinguistic studies of developing text production abilities in early and late childhood from several different perspectives. For example, Katzenberger (1994: 74–76) found a very high level of agreement between eight different judges, half experts (in educational psychology, literacy, and speech pathology) and half lay men and women asked to rank the narrative proficiency of texts produced by 25 preschoolers and five adults based on the same picture series. Similar findings emerged from the questionnaire addressed by Peterson and McCabe (1983) to seven judges, five experts and two lay people, who were asked to rank 288 personal experience narratives told by children from 4 to 10 years of age on a six-point scale. Relatedly, Malvern, Richards, Chipere and Duran (2004), working in a totally different framework, found that readers generally agreed very broadly in their overall evaluations of schoolchildren's written compositions. Ravid and Katzenberger (1999) also found very high, statistically very significant, levels of internal consistency between the global evaluations of members of a research team, specialists and graduate student majors in linguistics, literature, literacy, psychology, education, and speech pathology, who ranked 240 narrative and expository texts produced by Israeli schoolchildren, adolescents, and adults on a score from 1 to 5.

* This is a revised and expanded version of a paper presented at the Conference on New Directions in Cognitive Linguistics, University of Sussex, UK, October 2005. The authors are grateful to participants for their feedback, and we are indebted to the editors of this volume and to an anonymous reviewer for extremely valuable comments on the original draft. Responsibility for inadequacies rests with us alone.

On the other hand, in-depth probing of the motivations underlying these rankings led the authors to conclude that each of the seven team members had applied very different criteria in evaluating the texts as relatively good or poor.

Such studies provide empirical evidence for our claim that the notion of ‘text quality’ is elusive from two points of view. First, judges from different backgrounds and with different motivations vary widely in what they attend to, and what they consider important when reading a text – ranging from originality and literary style down to spelling and punctuation. Second, even experts seem to have difficulty in defining explicit, let alone precise, criteria that they use in evaluating texts.

In order to address this issue, and to fine-tune the idea of ‘text quality’, we examined texts produced by children, adolescents, and adults, native speakers of two typologically different languages, Californian English and Israeli Hebrew. In the present context, focus is on the essays – that is, texts constructed in the expository genre of discourse and in the written modality – that constituted part of a larger database of texts produced by participants in a large-scale cross-linguistic study on developing text construction abilities (Berman and Verhoeven 2002; Berman 2005).¹ We deliberately selected *expository discourse* as the locus for this study, to represent what Bruner (1986) terms the “logico-scientific paradigm” compared with the “narrative mode of thought”.

A major reason for deciding to focus on expository essays in our sample is that there is a rich literature on *narrative* discourse in cognitive science (e.g. Rumelhart 1975; Mandler 1987), in linguistics (e.g. Labov 1997; Longacre 1996) and in developmental psycholinguistics (e.g. Berman and Slobin 1994; Hickmann 2003). Such studies have specified criteria for evaluating different facets of narrative text construction, in terms of such notions as cohesion and coherence, a narrative schema, narrative evaluation, and/or maintaining and shifting reference. In contrast, far less (psycho-) linguistically motivated research is available on expository discourse in general, and on developing text construction abilities in this domain in particular. (A notable exception is Scinto’s (1986) analysis of school age writing, but even this fails to provide explicit criteria for what constitutes a ‘good’ essay beyond the notions of compactness and coherence.)

Besides, the principles underlying the organization of expository discourse are not only less clearly specified in the research literature than for narratives, but expository discourse in essence is less immediately accessible to such definition. This is because expository texts are by nature more abstract, since they focus on topics and issues, and express the unfolding of ideas, claims, and arguments that are logically interrelated with one another (Britton 1994; Katzenberger 2004; Mosenthal 1985). In this, they contrast markedly with narratives, which focus on people, their actions and motivations, and express the unfolding of events in a temporal framework, so that narrative content is to a large extent dictated by external events (occurrences in the real or fictive world) arranged in sequence, and constructing an effective narrative depends largely on pragmatic, communicative, and affective skill at ‘storytelling’ (Berman 1995; Reilly 1992). True, there is general agreement on the centrality of

1. The larger study elicited texts in seven different languages from schoolchildren, adolescents, and adults in four different age-groups, each of whom produced four texts – in two different genres (personal experience narratives versus expository discussions), and in two modalities (speech and writing).

narrative discourse in cognitive activities and in the way people interpret the world (Gerrig 1993; Turner 1996). Nonetheless, and perhaps for this very reason, expository discourse as a symbolic activity appears rather more interesting in the present context for examining the relation between language and thought. Besides, in contrast to narratives, rather than describing events that have or could have occurred, expository discourse *creates its own content*, so to speak. As Britton (1994) points out, the very function of expository texts is “to create a thematic structure in the reader’s (or hearer’s) mind”. As a result, expository texts show a very close connection between discourse structure and thematic content. Hence, “the structuring of a piece of expository discussion depends not only on how the flow of information is organized but also on the logical consistency and originality of the propositional content which it conveys” (Berman and Katzenberger 2004: 89). Moreover, while narrative discourse is essentially a way of performing or interpreting human reality, an expository text goes beyond this in order to reconstruct and verify reality by a dialectic relationship of analysis and synthesis (Georgakopoulou and Goutsos 2000). In consequence, expository discourse is constrained *from within*, and as such its construction imposes a heavy burden in both processing and ideation, one that presents even mature speaker-writers – let alone school-age children – with a major cognitive challenge.

2. Analyses and findings

To test the question of the relation between language and cognition in developing (expository) text construction, we examined the connection between quantitative and qualitative text measures applied to essays written by 160 children, adolescents, and adults – native speakers of Californian English and Israeli Hebrew – who had been asked to write an essay or composition discussing the topic of interpersonal conflict or ‘problems between people’.² The present study is thus in some ways comparable to the work of Malvern et al. (2004), who looked at the relationship between quantitative text variables and the quality of writing of (narrative) texts produced by schoolchildren of different ages (7, 11, and 14 years of age). Our analysis, however, involves rather different criteria for comparing text construction abilities along the following two dimensions: (1) *Local linguistic expression* – as measured by lexical usage and clause-level syntax; and (2) *Global discourse quality* as measured by structural well-formedness – in terms of the integration of top-down and bottom-up processes of text construction and of pragmatic appropriateness as reflecting individual rhetorical style. We selected these dimensions for analysis as respectively reflecting more directly linguistic, ‘item-based’ facets of text construction abilities (Berman 1986, 2004; Karmiloff-Smith 1986) and more general cognitive underpinnings

2. Subjects were given minimal “scaffolding”, beyond a brief background trigger of wordless video clips depicting (unresolved) situations of conflict between young people coupled with the instruction to write an essay or composition expressing their ideas on the topic of “problems between people”. They were also told explicitly “to discuss the topic” and *not* to tell a story in so doing. Details of elicitation procedures are provided in Berman and Verhoeven (2002), Berman and Katzenberger (2004), Berman and Nir-Sagiv (2007).

of this domain. That is, in order to treat local linguistic expression and global discourse organization as two separate facets of text construction abilities, we deliberately factored out the dimensions of vocabulary and grammar from our evaluation of global text quality, and, conversely, we disregarded overall text structure and content in evaluating local linguistic expression.

2.1 Local linguistic expression: Vocabulary and syntax

In the interests of brevity, and because analyses of the local level of linguistic expression in the same or similar samples as our own are documented elsewhere, we merely summarize our measures and key findings for this dimension. First, in the domain of vocabulary, we applied the following measures: (1) *Word length* – in both English and Hebrew, by number of syllables, in English by number of letters as well, comparing the number of polysyllabic words (three syllables or more) across the variables of age and text type (Nir-Sagiv 2005); (2) *Lexical density* – in both languages, measured by the proportion of open class items or content words (nouns, verbs, adjectives) out of the total words per text (Nir-Sagiv, Bar-Ilan and Berman 2008; Ravid 2004; Strömquist et al. 2002); (3) *Register* – in the sense of level of usage, specified as formal versus colloquial for English (Bar-Ilan and Berman 2007) and as high, neutral, or low for Hebrew (Nir-Sagiv, Sternau, Berman and Ravid 2008; Ravid and Berman 2009); and (4) *Noun abstractness* – measured by a condensed version of the 10-point scale developed by Ravid (2006), ranging nouns in terms of semantic content from concrete, imageable, and specific at one end to abstract, generic, and derivationally complex at the other (Berman and Nir-Sagiv 2007). Figure 1 summarizes the findings collapsed for distribution of vocabulary usage across these four measures in the expository essays written in both English and Hebrew, across the four age groups.

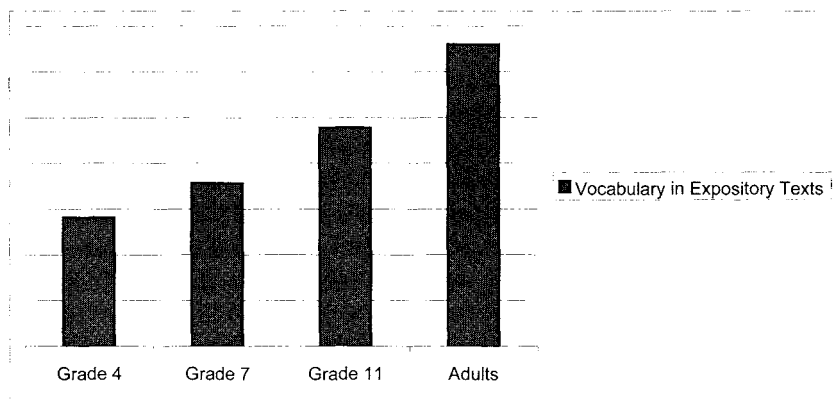


Figure 1. Overall patterning of lexical usage in English and Hebrew combined, by age [N = 40 in each age group]

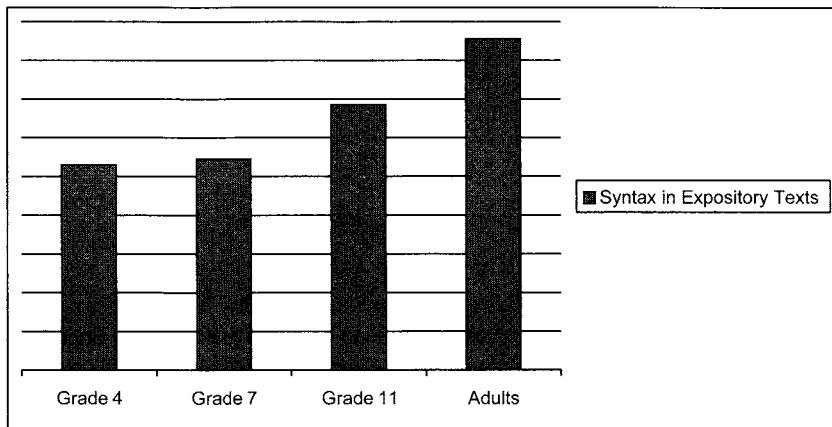


Figure 2. Overall patterning of syntactic complexity in English and Hebrew combined, by age [N = 40 in each age group]

As shown in Figure 1, the combined score for all the criteria of lexical usage (length, density, register, and nominal abstractness) taken together, for both English and Hebrew, revealed clear and significant differences by age and by genre. Across the board, younger children – grade schoolers in middle childhood (aged 9 to 10 years) and junior high pre-adolescents (aged 12 to 13 years) – score lower than high school adolescents and graduate student adults. And scores were higher in expository texts across the population.

In the domain of *syntax*, we applied the following measures: (1) *Clause length*, measured by mean number of words per clause, as an indicator of clause-internal syntactic density (Berman and Verhoeven 2002: 25–28; Chafe 1982; Hunt 1965; Ravid 2004: 343–346);³ (2) *Noun phrase complexity*, measured by a procedure designed for the larger cross-linguistic project, and involving three different criteria that included length in words, number and type of modifiers, and syntactic depth (Ravid and Berman in press); (3) *Relative clause constructions*, known to be a late developing usage in children’s oral narratives and characteristic of advanced level writing in different languages (Berman 1998; Dasinger and Toupin 1994; Loban 1976, Scott 1988), as corroborated by findings for the sample under consideration here (Berman and Nir-Sagiv 2009; Verhoeven et al. 2002: 152–153). Figure 2 depicts the overall patterning of these three syntactic measures taken together for both English and Hebrew, across the variable of age.

Figure 2 reveals a remarkably similar picture for syntactic patterning to what was observed for lexical usage across the variables of both age and genre. From the youngest age groups up, there were significant age-related differences along the different measures,

3. The “clause”, defined as a “unit that contains a unified predicate” (following Berman and Slobin 1994: 660–664), constitutes the basic unit of our text-based analyses across age-groups and text types in different languages.

there was a marked jump between the junior high and high school age groups, and more complex syntactic usage occurred in expository than in narrative texts.⁴

In sum, across a range of different measures of linguistic expression, in both lexicon and syntax and in both English and Hebrew, statistically significant variation emerged across *age*: older speaker-writers used lower frequency, more advanced vocabulary and more marked and complex syntactic constructions than the younger children, and there was a clear cut-off point at high school, with adolescents and adults clustering together in contrast to 9- and 12-year-olds. Other studies in the same framework further revealed a significant impact of *genre*: in different languages, including English and Hebrew, expository texts consistently elicit more sophisticated vocabulary, higher register usage, and more complex syntax than the narratives produced by the same subjects on the shared topic of interpersonal conflict (Berman 2008).

Moreover, while these general trends apply very similarly across the populations in both English and Hebrew, as is to be expected, *cross-linguistic differences* emerge as relevant to evaluation of local linguistic expression. For example, the historically determined 'Germanic-Romance' divide of English vocabulary into relatively colloquial compared with more formal, high register items is highly language-particular. Thus, while vocabulary likewise played a predominant role in the specification of Hebrew linguistic forms as of high, neutral, or low register, the historical antecedents of these terms (specifically, whether of classical Biblical or of post-Biblical, Mishnaic origin) could not be identified as denoting a high as against a more everyday level of usage respectively. And in the domain of syntax, as noted, use of passive voice and of nonfinite subordination is far more relevant to defining complex syntactic usage in English than in Hebrew. In Hebrew, in contrast, derivational morphology plays a major role in the lexicon-syntax interface, since it is used to mark valence relations on verbs (including reliance on middle voice constructions for downgrading of agency) as well as for constructing denominal and deverbal nominalized constructions in a particularly dense grammatical packaging of information (Berman 2004; Ravid and Cahana-Amitay 2005).

A fourth finding that applied consistently across age groups and languages was the fact that significant correlations emerged between the different criteria of local linguistic expression that we measured. We found this not only between the different areas of vocabulary usage (Nir-Sagiv, Bar-Ilan and Berman 2008) but also through calculations undertaken specifically for the present analysis, in the even stronger relationship that emerged between *lexicon* and *syntax*. Table 1 shows the correlations between each of the different criteria of vocabulary and between these and the syntactic constructions we measured for Hebrew and for English.

As can be seen from Table 1, the four measures of vocabulary correlate significantly with both clause and noun-phrase complexity, with correlations varying in magnitude

4. Two other measures of syntactic complexity – use of passive voice (Jisa et al. 2002; Reilly et al. 2005) and nonfinite subordination (Berman and Nir-Sagiv 2009; Kupersmitt 2006) – revealed closely similar trends in English, although largely irrelevant to Hebrew for reasons of typological differences between the two languages.

Table 1. Correlations between measures of vocabulary and of syntax, by age (grade schoolers [G] and junior high [J] versus high school students [H] and adults [A]) and by language [N = 40 per group]

Correlations			Clause length	Relative clauses	NP complexity
Spearman's rho					
G+J	Hebrew	Lexical Density	0.74**	-0.25	0.02
		Linguistic Register	0.33	0.12	0.18
		Semantic Abstractness	-0.23	0.01	0.49*
		Word Length	-0.25	0.05	0.32
	English	Lexical Density	0.52*	0.02	0.59*
		Linguistic Register	0.09	0.12	0.14
		Semantic Abstractness	0.09	-0.46	0.53**
		Word Length	-0.16	0.14	-0.34
H+A	Hebrew	Lexical Density	0.80**	0.23	0.45*
		Linguistic Register	0.14	0.20	0.55*
		Semantic Abstractness	0.45*	0.07	0.57**
		Word Length	0.15	0.16	0.10
	English	Lexical Density	0.74**	0.24	0.36
		Linguistic Register	0.34*	0.12	0.42*
		Semantic Abstractness	0.34	0.28	0.73**
		Word Length	-0.04	0.09	0.20

** Correlation is significant at the .01 level (2-tailed)

* Correlation is significant at the .05 level (2-tailed)

from medium to strong.⁵ The finding that vocabulary measures failed to correlate significantly with mean number of relative clauses per text can be attributed to the fact that relative clauses in general were rather rare across the sample, in marked contrast to broad-based representation of clauses and noun phrases. Overall, our findings for vocabulary-syntax correlations in expository texts produced from middle childhood across adolescence confirm the strong interconnection that has been observed between command of the lexicon and grammatical development in young pre-school children acquiring different languages (Caselli, Casadio and Bates 1999; Fenson et al. 1994; Marchman and Bates 1994; Marchman and Thal 2005) and it supports a psycholinguistic view of grammar and the lexicon as 'inseparable' (Bates and Goodman 1997).

2.2 Global text quality

As a next step, we went beyond these more item-based and closely interrelated local measures of language usage to consider what we term 'global text quality'. Our aim here was to provide explicit, and empirically well motivated criteria of overall text construction

5. Bolded figures in Table 1 also indicate medium-sized, though statistically non-significant correlations.

and the organization of information across an entire piece of discourse, independently of specific features of lexicon and grammar. This involved, first, specifying explicitly what constitutes structural well-formedness and, second, the idea that proficient text construction requires speaker-writers to go ‘*beyond well-formedness*’. We propose that any piece of discourse is defined by genre-dependent constraints of ‘*structural well-formedness*’. In expository texts, this involves fleshing out the top-down discourse topic (here, the subject of interpersonal conflict or ‘problems between people’) with bottom-up categories that elaborate on, illustrate, and provide evidence for the top-down generalities. That is, well formed text construction requires an interplay between top-down and bottom-up processing – with different principles of organization in different genres, from top-down generalizations to specific categories in expository discussion, as compared with the shift from bottom-up events to a global narrative schema in stories. Going beyond well-formedness involves the ability to diverge from genre-canonic structure and content, by reference to narrative-like personal experiences or culturally shared past events in expository texts, analogous to making expository-like generalizations in narratives (Berman and Nir-Sagiv 2004; Kupersmitt 2006).⁶ It also means being able to enhance a piece of discourse by meta-textual and meta-cognitive allusion or skillful use of rhetorical devices. These latter abilities characterize a text that is not only ‘well formed’, but a ‘good’ piece of discourse, in the case at point here, an original, interesting, and illuminating discussion of an abstract topic.

Our analysis of overall ‘text quality’ thus takes into account both discourse-structural competence and text production performance in order to specify what is involved in fully proficient text construction.⁷ As noted elsewhere, for development of oral narratives, the line between competence and performance is “not only flexible and fuzzy, but ... bi-directional. Knowledge of linguistic forms and narrative structure clearly underlies the ability to tell a story; but the act of (story hearing and) storytelling impinges on this knowledge and affects it across the developmental history of each individual in becoming a proficient interpreter and teller of stories” (Berman 1995: 308). In the present analysis, achieving structural well-formedness requires that, just as children learn to integrate events within an overall narrative schema in telling a story, so in constructing an expository text, they need to constrain their generalizations by categorical specificities.

This view of text construction enabled us to define the notion of ‘global text quality’ for expository as well as other types of texts on the basis of three distinct but related dimensions of analysis (again, as analyzed quantitatively, motivated, and illustrated extensively in Berman and Nir-Sagiv 2007). The first relates to the level of *cognitive representation* – from minimal reliance on only bottom-up (narrative events) or top-down (expository generalizations) to a creative synthesis, requiring what Karmiloff-Smith (1992) terms “integrated

6. This alternation between “genre-canonic” and “genre-extraneous” elements of a text as an important facet of developing text construction abilities and of discourse proficiency in general is described and motivated in detail in Berman and Nir-Sagiv (2007).

7. Our use of these terms departs from Chomsky’s (1965: 3–15) distinction between *competence* as underlying knowledge of (linguistic) structure compared with *performance* as use of this knowledge in behavioral terms of language comprehension and production.

Table 2. Four levels of text construction abilities defined for narrative and expository discourse along the three dimensions of level of representation, structural well-formedness, and discursive effectiveness*

Level → Dimension ↓	I	II	III	IV
Representation and Cognitive Processing	<i>Minimal Representation:</i> Relying only on Top-Down generalizations	<i>Partial Extension:</i> Initial integration of Top-Down / Bottom-Up categories	<i>Well-Formedness:</i> Full integration of Top-Down and Bottom-Up	<i>Beyond Well-Formedness:</i> Creative synthesis of parts into whole
Structure and Content	Only basic components	Initial reliance on genre-typical features	Overt categorial structure	Genre-external material
Discursive Features	Detached units	Initial anchoring	Relating beginning and ending; explicit departures from genre-canonic	Meta-cognitive inter-textual and / or meta-textual

* Entailment relations obtain between the four levels, such that Level II precludes Level I, and Level IV implies having achieved Level III

re-representation”, here of top-down and bottom-up processing in the course of ongoing text construction. The second concerns level of *discourse structure* and content – from only core or obligatory to optionally elaborative elements. In expository discourse, these take the form of ‘move-on’ statements or ‘core propositions’ that are elaborated by illustrative or delimiting ‘satellite’ elements (Britton 1994; Matthiessen and Thompson 1988). The third dimension of text construction involves features of *informative content* – including departures from genre-canonic features (for example, by including narrative-like incidents from personal or past history in an expository text) and reliance on rhetorical strategies such as parallelism or inter-textual and meta-textual commentary.

Against this background, we delineated *four levels* for ranking global level text construction. As summarized in Table 2, each level was specified along three interrelated dimensions, as follows: (1) Level of cognitive representation and processing was the basis for what we specified as the *defining rank* of a given text: from I – *minimal* reliance on only top-down generalities via II – *partial extension* and initial integration of top-down and bottom-up categories, to III – well formed *full integration* of top-down and bottom-up elements of discourse, culminating in IV – a *creative synthesis* of the parts into a whole; (2) Discourse structure and content yielded what we termed the *critical properties* of text construction, ranging from I – only *basic elements* of the genre such as narrative events or detached expository generalities via II – initial reliance on *genre-typical* features such as a narrative orientation or minimal specificities in expository texts, to III – a fully explicit *schematic discourse structure* in narratives and categorial specification in expository texts, and on to IV – overt reliance on *genre-extraneous material*; and (3) *characteristic features* of discourse, ranging from I – *detached units* not anchored in any other text-internal content, to II – *initial anchoring* by reference to the elicitation procedures (in this instance, the video shown at the beginning), via III – *fully anchored* texts that relate the opening

and closing, the introduction and conclusion, and on to IV – sophisticated *meta-textual manipulation* of text content, by inter-textual allusions, deviations from genre-canonic features, and/or self-conscious reliance on rhetorical devices.

These four levels of text construction are each illustrated by an expository text written by native speakers of Californian English in (1) to (4) below, with participant ID indicated for language (e for English), grade level (G, J, H, A, see Table 1), and serial number in that age group (01 to 20).

(1) LEVEL I: Composition Written by Fourth Grade Girl [eG16]

I do not think fighting is good. You do not make friends that way. If you do not fight, you can have many many friends. But when you fight, you can hurt the person's feelings you are fighting with. You should always be nice and respectful to other people. And if you are not nice, you will end up not having any friends. That is why you should not fight.

This 9-year-old's text makes good use of local-level vocabulary and grammar, but in global text construction, it fails to go beyond what we termed Level I, since it contains only 'minimal' top-down components of expository structure in the form of detached generalized propositions, not anchored in any specific reality. In contrast, the 12-year-old boy's text in (2) was ranked at Level II since, although it remains in the realm of generalizations, there is some initial attempt at sub-categorization to the effect that different people handle problems in different ways, and that there are both small and larger conflicts and problems, and that these take place in different places.

(2) LEVEL II: Composition Written by Seventh Grade Boy [eJ05]

I think there are many problems and conflicts in the world. I also think different people handle these problems in different ways. Some people make little problems out to be big conflicts. The world has many huge problems that need to be dealt with a lot quicker than some people's little problems. Some problems can lead to many bad conflicts, which happens a lot at schools, on the street, and many other places. Little problems can be easily set aside, while big problems might take thinking and some action. Different people can lead to many problems and differences in opinions also lead to many problems. I think if you are a good person, you can overcome most problems in life.

The Level II text in (2) represents what we termed *partial extension* of core elements through addition of informational components and initial combining of top-down generalities and bottom-up specific categories. Nonetheless, like the text in (1), this remains within the realm of generalizations in the timeless present. The text in (2) also contains irrealis, modally formulated projections which, in contrast to (1), include epistemic reference to what *might* happen and are not confined to deictic commenting on what *should* happen (Reilly et al. 2002). The text written by a 17-year-old in (3) differs from both these quite markedly.

(3) LEVEL III: Composition Written by Eleventh Grade Boy [eH01]

Conflict is a large problem particularly in High School, although it never goes away. High School is a major focal point of conflict because of the extreme amount of new tension that students are confronted with. Coming from a sheltered environment with the close supervision and intervention of parents and teachers students are thrust into realization of the so called 'real world' where you must now make choices and resolve problems on your own. While you

are never really on your own, this new freedom can give the overwhelming feeling of distancing yourself from your parents' control. Students are exposed to many new people and begin to form social cliques or groups. These groups not only follow racial and ethnic lines but also the class bracket that they are placed in such as advanced or remedial. This can have an impact on people because of the exposure or lack of it or jealousy and envy. Peer pressure is one of the main causes of conflict which never goes away but that students have a hard time learning to cope with. While conflict is not a necessarily bad thing, it does help prepare people for the real world which is full of conflict and problems.

The text in (3) was rated as representing fully *well formed* discourse structure, since it integrates its top-down generalizations about the topic of conflict with bottom-up specification about how and why conflict takes shape in the specific context of high school as 'a major focal point of conflict' and as representative of the so-called 'real world'. Sub-categorization is expressed in terms of different social groups and factors such as jealousy and peer pressure, and the text is rounded off by an explicit concluding statement. On the other hand, this text was not rated as going 'beyond well-formedness', since it remains entirely genre-canonic in content and expression, with no illustrative instances anchored in past time and no reference to personal experience or individual knowledge. These features are, however, clearly evident in the text in (4), written by a 16-year-old girl in the same grade. This was rated at the highest of the four levels of text construction abilities, since it manifests transitions from generalizations to specific examples and from a generic, distanced discourse stance to a more personalized, and immediately involved one – as indicated by the underlined material.

(4) LEVEL IV: Composition Written by Eleventh Grade Girl [eH04]

Conflict is opposing ideas or stances between two or more people. In many ways it is a necessary part of life. On the other hand, it can cause disruption and chaos in the relationships of those involved. When people have a difference of opinion, a conflict is usually the result. This is a good way for those differences to be put aside. For example, I recently started swimming under a new coach. I did not like him at all – his coaching styles, the swim sets he assigned, or his overall attitude. One day after practice, I approached him and told him what bothered me about him, and that it was affecting my attitude about swimming. Since then we have gotten along much better, and I have a new appreciation for his coaching style, because he further explained it to me. In that way, conflict can be a good thing. The results that were achieved were better than the situation beforehand. On the contrary, conflict can ruin a friendship. My friend was very close friends with two other girls. They had an argument over priorities. The other girls would make plans and then break them at the last minute. Since then my friend has not spoken with them. This is a situation in which conflict was a bad thing. If the conflict cannot be resolved, then the relationship will suffer. In my case, I avoid conflict at all costs. It sometimes gets to a point where I void my opinion or hide the truth in order to prevent a conflict. On one hand, I very seldom argue with people. On the other hand, my ideas may go unheard, or a friendship is based on a fake foundation. There is a happy medium somewhere though. Hopefully some day I will realize when a conflict is necessary and use it intelligently, not as a fight but as a discussion to solve a common problem.

The text in (4) is not only structurally well formed in proceeding from top-down topic-based generalizations to specific sub-categories and back again, it goes beyond well-formedness

by providing narrative-like illustrations from personal experiences in the past (her own with her swimming coach and her friend's with other girls) and by expressing the writer's meta-cognitive analytical inferences about herself in relation to the topic of conflict. This is an example of what we termed 'effective text construction', with clear shifts between the expository and narrative modes of discourse and rhetorical expressiveness achieved, *inter alia*, by explicit marking of logical relations between parts of the text such as 'on the contrary', 'on one hand ... on the other hand'. That is, this Level IV text is both structurally well formed and contains fully integrated information, by transitions between general statements and specific instances, between expository- and narrative-like segments, and between a distanced, impersonal and generalized discourse stance to subjective, personalized commentary on the topic of discussion. These abilities, as shown below, proved beyond what younger children were capable of achieving.

The criteria defined for the four levels of global text quality, as summed up in Table 2 and illustrated by the texts in (1) through to (4), were applied to the 320 texts in our sample – both the narrative and the expository text written by each of 80 subjects in the four age groups in both English and Hebrew. The two genres are represented here, to highlight the developmental contrast between expository and narrative text construction. Findings for each of the four levels of global text quality breakdown are given in

Table 3a. Breakdown of English texts (in raw figures) into four levels of discourse construction abilities, by genre and age level [N = 20 per group]

Group	Narrative					Expository				
	0*	I	II	III	IV	0	I	II**	III	IV
G	1	0	14	5	0	1	14	4	1	0
J	2	0	5	9	4	0	1	16	1	2
H	1	0	0	9	10	0	0	7	9	4
A	0	0	0	5	15	0	0	0	11	9

* A 'zero' was assigned in cases where participants failed to construct the elicited text type, but instead produced a text in a completely different sub-genre, for example a description of the video clip that was used for elicitation.

** Texts produced by two subjects in Group J and four subjects in Group H were coded as between level (II-III). For purposes of tabulation, these subjects were entered at Level II, but they were given half-points for statistical analysis.

Table 3b. Breakdown of Hebrew texts (in raw figures) into four levels of discourse construction abilities, by genre and age level [N = 20 per group]

Group	Narrative					Expository				
	0	I	II	III	IV	0	I	II	III	IV
G	2	2	9	7	0	4	7	9	0	0
J	3	0	3	13	1	0	5	13	2	4
H	0	0	2	12	6	0	1	7	11	1
A	0	0	0	6	14	0	0	3	11	6

Tables 3a (for English) and 3b (for Hebrew), with figures in bold representing the most favored level for each age group in each of the two discourse genres.

These figures demonstrate that the criteria we devised for ranking global text quality can be considered both valid and reliable. First, its validity is demonstrated by the clearly consistent and statistically significant differences between levels of global text quality in the two genres (for details, see Berman and Nir-Sagiv 2007). Across age groups, narrative text production emerged as 'one step ahead' of expository texts, in the following sense. In *narratives*, the youngest subjects in our study – 9- to 10-year-old 4th graders – were nearly all beyond the first level defined here (and several reached Level III), while older subjects, all the adults and most of the high school students, demonstrated fully proficient, mature narrative abilities. In contrast, both across and within age groups, subjects *scored lower* in expository than in narrative text construction, with statistically significant differences across the variables of age as well as genre. This finding is supported by analyses of text openings and closings in the English, Spanish, and Swedish texts elicited in the cross-linguistic project (Tolchinsky, Johannssen and Zamora 2002) and of text openings in English and Hebrew (Berman and Katzenberger 2004). These developmental trends in different studies on related databases provide strong empirical support for the independent analyses we constructed for evaluating global text quality. The reliability of the global-level analyses presented here is further confirmed by the very marked similarity between findings for the two languages, Californian English and Israeli Hebrew, in texts elicited from two different populations. There was, moreover, a high rate of inter-judge reliability for both languages ($\kappa = 0.93$ for English, and 0.94 for Hebrew).⁸ Taken together, these findings demonstrate that the characterization of global text quality delineated in this section is both quantifiably reliable and conceptually robust.

2.3 Interrelation between local and global levels of text construction

We come now to the key question of this study: What, if any, is the relationship between local linguistic expression and global text construction? As noted earlier, the only relevant research we know of that directly addresses this question is Malvern et al.'s (2004) detailed and extensive research project covering large populations of British schoolchildren.⁹ In attempting to specify the relations between local linguistic expression and global text

8. All analyses were conducted on all texts by the authors (a native speaker of English and of Hebrew respectively). Thanks to Irit Katzenberger and Judy Kupersmitt for conducting the reliability checks.

9. Careful psycholinguistic criteria were also applied in the framework of the so-called 'nun study' on changes in linguistic ability in adulthood, focused on processes of aging (e.g. Kemper et al. 2001; Mitzner and Kemper 2003). Thus, the Kemper et al study takes into account both grammatical complexity and idea density (or propositional density), defined as "how much information can be packed into a sentence, relative to the number of words" (2001:228), to show that "idea density is less affected by normal aging processes than grammatical complexity (whereas) Alzheimer's disease leads to a rapid decline in idea density" (op.cit). Yet these studies fail to take into account the issue of global text quality in their analyses of autobiographical accounts.

quality in expository texts, our study both relies on and goes beyond such prior research in several ways.

In evaluating levels of linguistic expression, we take into account item-based facets of lexical *and* syntactic complexity and the correlation between them (Section 2.1), whereas Malvern and his associates considered different (also quantifiable) word-level variables (length, rarity, and spelling) supplemented only by measures of text length and T(erminable)-unit length (pp. 152–160). Besides, our analysis is more explicit and more inclusive than other developmentally motivated evaluations of expository text construction (Katzenberger 2004; Scinto 1986). More importantly, the model we propose diverges considerably from accepted educational assessments of ‘text quality’ (for example, the Qualifications and Curriculum Authority [QCA] as stipulated in the English National Curriculum used in the Malvern et al. study, or the United States’ National Assessment of Educational Progress NAEP). These and related studies, as noted earlier, typically reveal high inter-judge reliability, but they fail to tease apart distinct dimensions of what is involved in how ‘good’ a text is deemed to be. They tend to involve a mixture of criteria, ranging from overall text structure, thematic content, or originality, at one extreme, to language-specific features like vocabulary, grammar, spelling, and punctuation, at the other. Against this background, we decided in principle to avoid overlapping criteria in our evaluation of overall text quality as against local linguistic expression. As a result, our analysis of global text quality relies heavily on facets of cognitive processing defined in terms of top-down and bottom-up components of discourse, while our analysis of linguistic expression is confined to lexical items and syntactic constructions.¹⁰

In order to address the key question of concern to this study, we applied Spearman-rho correlations, controlling for age, between values on all the local measures of vocabulary and syntax (Section 2.1), on the one hand, and scores on global text quality (Section 2.2), on the other, supplemented by additional scores for text length measured in both words and clauses.

Table 4 shows that the correlations we tested for revealed *no* significant connection between local linguistic expression and global text quality – as defined by us – with the exception of overall *text length* – measured by mean number of both words and clauses per text. The fact that text length correlated quite highly with global text quality in each of the four age groups provides further support for our proposed model of overall discourse construction, since this finding is consistent with trends revealed by other studies of developing text construction (as further discussed below). On the other hand, local linguistic expression across varied measures of vocabulary usage (word length, lexical density, high or low register of usage, and nominal abstractness) and of syntactic complexity (clause length, relative clause usage, and noun phrase complexity) failed to reveal any significant correlation with global text quality (Table 4). In other words, the most conclusive finding from our analysis is for a relative *dissociation* between local versus global facets of text construction.

10. An interesting mediator between these two levels of text construction is provided by clause-linkage as a device for achieving textual connectivity (Berman 1998) by means of what we elsewhere term “clause packaging” (Berman and Nir-Sagiv 2009). Preliminary analysis suggests that this is a particularly fruitful domain for detecting cross-linguistic differences, of a kind not dealt with in the present context.

Table 4. Correlations between local linguistic measures and global text level for English and Hebrew combined, by age [N = 40 per group]

Spearman's rho	Global Text Quality			
	Grade 4	Grade 7	Grade 11	Adults
Text Length in Words	.553(**)	.727(**)	.518(**)	.561(**)
Text Length in Clauses	.487(**)	.664(**)	.496(**)	.514(**)
Lexical Density	-.197	-.164	.030	-.103
Linguistic Register	-.091	-.115	-.037	-.168
Word Length	-.014	-.156	-.099	-.244
Semantic Abstractness	.185	-.113	-.030	-.104
Words per Clause	.016	.085	.035	-.121
Relative Clauses	.026	-.048	.194	-.063
Noun Phrase Complexity	.102	-.109	.003	-.085

** Correlation is significant at the .01 level (2-tailed)

3. Discussion

This study has attempted to throw light on the elusive notion of 'text quality' by an innovative analytic heuristic aimed at teasing apart local and global facets of text quality. By precisely defining a range of different factors involved in *local linguistic expression*, we were able to uncover a strong correlation both between different measures of lexical usage on the one hand, and between lexicon and syntax on the other. This, as noted earlier, provides impressive evidence for the inseparability of lexicon and syntax in a usage-based linguistic approach and is in itself a noteworthy outcome of the present study: it sheds light on an issue that, to the best of our knowledge, has not been addressed in prior developmental research along such a broad and innovative range of quantifiable criteria. Overall, the categories we adopted for assessing local linguistic expression reveal 'text-embedded language usage' as a complicated, yet clearly measurable facet of text construction.

On the other hand, contrary to our initial assumption, we found a clear dissociation across the population in local and global dimensions of text quality. Nonetheless, *both* showed marked and significant changes as a function of age, with a major cut-off point between junior high 12- to 13-year-olds and high school adolescents. In other words, these two critical facets of text construction abilities both develop and change in tandem as a function of increased maturity and the extended experience with literacy-based activities required in order to write an expository essay on an abstract topic such as interpersonal conflict (Berman and Katzenberger 2004; Ravid 2004).

The only significant correlation between the measures we applied of local linguistic expression and global text quality was yielded by the variable of *text length*, as measured both by number of words and number of clauses (Table 4). This is consistent with the findings of the cross-linguistic project from which this study derives, which revealed a significant effect of age on text length – measured by number of words per text – across all seven languages and in all four text types (narrative and expository, spoken and written), "with a major difference in mean scores between the junior high and high school

students suggesting that the gross amount of text produced increases sharply with age” (Berman and Verhoeven 2002: 23). The factor of age also had a significant effect on the mean length of the oral narratives produced by preschoolers, schoolchildren, and adults in five different languages (Berman and Slobin 1994: 31), although the authors note there that “some adults across the languages produce shorter, more condensed versions than certain of the children” and they caution against identifying the ‘goodness’ of a narrative with its relative length.

Interestingly, the variable of text length was the best predictor of text level in the Malvern and Richards study as well (along with the major variable of their study, vocabulary diversity).¹¹ In fact, text length appears to be the only domain, apart from general developmental findings for more advanced lexical usage with age, that is clearly consistent with their analysis of nearly a thousand texts written by schoolchildren who were asked to write ‘a narrative composition’ beginning with the words “*The gate was always locked but on that day someone had left it open*” (Malvern et al. 2004: 154). Thus, the English researchers found a highly significant correlation between several different language measures and what they term ‘text level’ (ibid.: 153–154). But this disparity can be attributed to various differences in methodology, as noted throughout our study. These include the fact that, most importantly, we adopted totally different criteria from their schooling-based ratings of overall text quality in evaluating expository discourse compared with make-believe narratives respectively. Moreover, the analyses differ in that ours controlled for age and hence involved relatively small, and possibly more homogeneous groups of subjects at each developmental level.¹²

On the other hand, we do not consider the difference in text type (expository rather than narrative) to be a major factor, since we assume that a similar dissociation between local linguistic expression (vocabulary and syntax) and overall text quality will be found in narrative text construction too. However, narratives might demonstrate greater compatibility between local and global facets of text construction than expository discourse. This is suggested by the *cross-genre* comparisons afforded by our study, where the same person produced both a personal experience account and an expository discussion on the shared topic of interpersonal conflict. When these two text types are compared for local linguistic expression, earlier studies in the cross-linguistic project show that from as early as 4th grade (the youngest group in our sample), children make use of higher level, more advanced or ‘school-like’ vocabulary and more complex syntax in expository compared with narrative texts (see Berman and Nir-Sagiv 2007; Nir-Sagiv, Bar-Ilan and Berman

11. Our analyses discounted the factor of lexical diversity – ratio of different words to total words in a text – since the procedure devised by Malvern, Richards, and their associates at Reading (so-called VOCD) relates to word types in the sense of different word forms, rather than of different lexemes, whereas the latter constitute a more valid criterion for languages which have a highly rich morphology, like Hebrew, particularly when evaluating text-embedded vocabulary use across adolescence and into adulthood.

12. We would like to thank David Malvern and Brian Richards of Reading University for their generous collegial input and advice on statistical analyses. Our decision to conduct analyses by age (rather than, as they did, across the population with age as a mediating factor) was motivated by the fact that our study is focused directly on two different facets of text construction – local and global.

2008 – for English; Jisa 2004 – for French; Ravid 2004 – for Hebrew; and Tolchinsky and Rosado 2005 – for Spanish). Yet our current analysis shows that in command of global text structure, expository text construction lags significantly behind narrative production within and across age groups (Tables 3a, 3b).

This discrepancy can be attributed to a special kind of developmentally anchored dissociation of abilities in the domain of monologic text construction, which appears particularly marked in non-narrative genres. Children may be able to use ‘book language’ (Blank 2003) from relatively early in their schooling, by middle childhood, but the cognitive demands of constructing a well formed and coherent piece of expository discourse are met only around adolescence. This is shown by our breakdown into four levels of expository text construction in contrast to the narrative texts produced by the same subjects in English and in Hebrew (Tables 3a and 3b). While all the adults in both languages achieve structural well-formedness (Levels III and IV) in both types of texts, and 16- to 17-year-olds do so in the majority of cases (in narratives to a larger extent than in expository texts), younger children aged 12 to 13 years do so most of the time in narratives, but rarely in expository texts. This suggests that the cognitively demanding task of writing an expository essay on an abstract topic of discussion – and for the two younger groups of subjects, grade school and junior high students, producing a written account of a personal experience – requires quite formidable intellectual resources and considerable mental effort.

The fact that we found local level of language use to be largely dissociated from overall text quality goes against our initial assumption, but it is in line with the common-sense perception that ‘more is not necessarily better’. In our sample, there were several subjects in both languages who scored high on vocabulary and syntax, but reached only Level I or II in global text quality. And conversely, many of the subjects who scored high on text quality achieved only average and sometimes even poor scores on linguistic expression. That is, a text may make use of an elaborate and complicated style of expression and yet be not be a ‘good’ piece of discourse, whereas a short text written in ‘plain English’ may rate high on overall quality. Such disparities between performance on local linguistic criteria and on global text quality caution against equating complex, high-level linguistic expression with proficient text construction. And it suggests that *individual style* is a determining factor in how people realize the task of ‘thinking for writing’ (following Slobin 1996, 2003).

In sum, it appears that the cognitive demands involved in integrating top-down and bottom-up processes in the course of (even off-line, written, let alone online spoken) text construction – while necessarily making use of language – can and perhaps should be evaluated independently of linguistic knowledge, at least in the case of non-expert speaker-writers like the participants in our study. The dissociation between level of language usage and overall text construction revealed by the study leaves cognitive science with yet another query: the interrelation between cognitive and linguistic factors in developing discourse abilities.

References

- Bar-Ilan, L. & Berman R. A. 2007. The Latinate-Germanic divide and register differentiation in developing text production across adolescence. *Linguistics* 45 (1): 1–36.
- Bates, E. & Goodman, J. 1997. On the inseparability of grammar and the lexicon: Evidence from acquisition, aphasia, and realtime processing. *Language and Cognitive Processes* 12 (5/6): 507–586. [Special Issue on the Lexicon].
- Berman, R. A. 1986. A step-by-step model of language learning. In *Stage and Structure: Re-opening the Debate*, I. Levin (ed.), 191–219. Norwood, N.J.: Ablex.
- Berman, R. A. 1995. Narrative competence and storytelling performance: How children tell stories in different contexts. *Journal of Narrative and Life History* 5 (4): 285–313.
- Berman, R. A. 1998. Typological perspectives on connectivity. In *Issues in the Theory of Language Acquisition*, N. Dittmar & Z. Penner (eds), 203–224. Bern: Peter Lang.
- Berman, R. A. 2004. Between emergence and mastery: The long developmental route of language acquisition. In *Language Development across Childhood and Adolescence*, R. A. Berman (ed.), 9–34. [Trends in Language Acquisition Research 3]. Amsterdam: John Benjamins.
- Berman, R. A. 2005. Introduction: Developing discourse stance in different text types and languages. *Journal of Pragmatics* 37 (2): 105–124. [Special Issue on Developing Discourse Stance across Adolescence].
- Berman, R. A. 2008. The psycholinguistics of developing text construction. *Journal of Child Language* 35: 735–771.
- Berman, R. A. & Katzenberger, I. 2004. Form and function in introducing narrative and expository texts: A developmental perspective. *Discourse Processes* 38: 57–94.
- Berman, R. A. & Nir-Sagiv, B. 2004. Linguistic indicators of inter-genre differentiation in later language development. *Journal of Child Language* 31: 339–380.
- Berman, R. A. & Nir-Sagiv, B. 2007. Comparing narrative and expository text construction across adolescence: A developmental paradox. *Discourse Processes* 43 (2): 79–120.
- Berman, R. A. & Nir-Sagiv, B. 2009. Clause-packaging in narratives: A crosslinguistic developmental study. In *Crosslinguistic Approaches to the Psychology of Language: Research in the Tradition of Dan I. Slobin*, J. Guo, E. Lieven, S. Ervin-Tripp, N. Budwig, S. Özçalışkan & K. Nakamura (eds), 149–162. Mahwah, NJ: Lawrence Erlbaum.
- Berman, R. A. & Slobin, D. I. 1994. *Relating Events in Narrative: A Crosslinguistic Developmental Study*. Hillsdale, NJ: Lawrence Erlbaum.
- Berman, R. A. & Verhoeven, L. 2002. Developing text production abilities across languages and age-groups. *Written Language and Literacy* 5 (1): 1–44.
- Blank, M. 2003. Classroom discourse: A key to literacy. In *Speaking, Reading, and Writing in Children with Language Learning Disabilities*, K. G. Butler & E. R. Silliman (eds), 151–174. Mahwah, NJ: Lawrence Erlbaum.
- Britton, B. K. 1994. Understanding expository text: Building mental structure to induce insights. In *Handbook of Psycholinguistics*, M. A. Gernsbacher (ed.), 641–674. New York: Academic Press.
- Bruner, J. 1986. *Actual Minds, Possible Worlds*. Cambridge, MA: Harvard University Press.
- Caselli, M. C., Casadio, P. & Bates, E. 1999. A comparison of the transition from first words to grammar in English and Italian. *Journal of Child Language* 26: 69–111.
- Chafe, W. L. 1982. Integration and involvement in speaking, writing, and oral literature. In *Spoken and Written Language: Exploring Orality and Literacy*, D. Tannen (ed.), 83–112. Norwood, NJ: Ablex.
- Chomsky, N. A. 1965. *Aspects of the Theory of Syntax*. Cambridge, MA: M.I.T. Press.
- Dasinger, L. & Toupin, C. 1994. The development of relative clause functions in narrative. In *Relating Events in Narrative: A Crosslinguistic Developmental Study*, R. A. Berman & D. I. Slobin (eds), 457–414. Hillsdale, NJ: Lawrence Erlbaum.

- Fenson, L., Dale, P. S., Reznick, J. S., Bates, E., Thal, D. J. & Petrick, S. J. 1994. Variability in early communicative development. *Monographs of the Society for Research in Child Development* 59 (5): 1–185.
- Georgakopoulou, A. & Goutsos, D. 2000. Mapping the world of discourse. The narrative vs. non-narrative distinction. *Semiotica* 131 (1/2): 113–141.
- Gerrig, R. 1993. *Experiencing Narrative Worlds: On the Psychological Activities of Reading*. New Haven: Yale University Press.
- Hickmann, M. 2003. *Children's Discourse: Person, Space, and Time across Languages*. Cambridge: Cambridge University Press.
- Hunt, K. W. 1965. *Grammatical Structures Written at Three Grade Levels* (Research Report No. 3). Champaign, Ill: National Council of Teachers of English.
- Karmiloff-Smith, A. 1986. Stage/structure versus phase/process in modeling linguistic and cognitive development. In *Stage and Structure: Re-opening the Debate*, I. Levin (ed.), 164–190. Norwood, NJ: Ablex.
- Karmiloff-Smith, A. 1992. *Beyond Modularity: A Developmental Perspective on Cognitive Science*. Cambridge, MA: M.I.T. Press.
- Katzenberger, I. 1994. *Cognitive, Linguistic, and Developmental Factors in the Narration of Picture-Series*. PhD Dissertation, Tel Aviv University [in Hebrew].
- Katzenberger, I. 2004. The super structure of written expository texts: A developmental perspective. In *Perspectives on Language and Language Development*, D. Ravid & H. Bat-Zeev Shyldkrot (eds), 327–336. Dordrecht: Kluwer.
- Kemper, S., Greiner, L., Marquis, J. G., Prenovost, K. & Mitzner, T. L. 2001. Language decline across the life span: Findings from the nun study. *Psychology and Aging* 16: 227–239.
- Kupersmitt, J. 2006. *Temporality in Texts: A Crosslinguistic Developmental Study of Form-Function Relations in Narrative and Expository Discourse*. PhD Dissertation, Bar Ilan University.
- Labov, W. 1997. Some further steps in narrative analysis. *Journal of Narrative and Life History* 7: 395–415.
- Loban, W. 1976. *Language Development: Kindergarten through Grade Twelve*. (Research Report No. 18). Champaign, Ill: National Council of Teachers of English.
- Longacre, R. E. 1996. *The Grammar of Discourse*. New York: Plenum.
- Jisa, H. 2004. Growing into academic French. In *Later Language Development: Typological and Psycholinguistic Perspectives*, R. A. Berman (ed.), 135–161. [Trends in Language Acquisition Research 3]. Amsterdam: John Benjamins.
- Jisa, H., Reilly, J., Verhoeven, L., Baruch, E. & Rosado, E. 2002. Passive voice constructions in written texts. *Written Languages and Literacy* 5: 163–182.
- Malvern, D. D., Richards, B. J., Chipere, N. & Durán, P. 2004. *Lexical Diversity and Language Development: Quantification and Assessment*. Hampshire: Palgrave Macmillan.
- Mandler, J. M. 1987. On the psychological reality of story structure. *Discourse Processes* 10: 1–30.
- Marchman, V. & Bates, E. 1994. Continuity in lexical and morphological development: A test of the critical mass hypothesis. *Journal of Child Language* 21: 339–366.
- Marchman, V. & Thal, D. 2005. Words and grammar. In *Beyond Nature-Nurture: Essays in Honor of Elizabeth Bates*, M. Tomasello & D. I. Slobin (eds), 139–164. Mahwah, NJ: Lawrence Erlbaum.
- Matthiessen, C. & Thompson, S. A. 1988. The structure of discourse and 'subordination'. In *Clause Combining in Grammar and Discourse*, J. Haiman & S. A. Thompson (eds), 275–330. Amsterdam: John Benjamins.
- Mitzner, T.L. & Kemper, S. 2003. Oral and written language in late adulthood: Findings from the nun study. *Experimental Aging Research* 29: 457–474.
- Mosenthal, P. B. 1985. Defining the expository discourse continuum. *Poetics* 14: 387–414.
- Nir-Sagiv, B. 2005. Crosslinguistic and developmental perspective on word length as criterion for vocabulary complexity. Paper presented at the International Association of the Study of Child Language (IASCL), Berlin.

- Nir-Sagiv, B., Bar-Ilan, L. & Berman, R. A. 2008. Vocabulary development across adolescence: Text-based analyses. In *Language Education in Israel: Papers in Honor of Elite Olshtain*, I. Kupferberg & A. Stavans (eds), 47–74. Jerusalem: Magnes Press.
- Nir-Sagiv, B., Sternau, M., Berman, R. A., & Ravid, D. 2008. Register differentiation as a feature distinguishing school-age usage across the variables of genres (narrative/expository) and modality (written/spoken). *Israel Journal of Language and Literacy* 1: 71–103. [in Hebrew].
- Peterson, C. & McCabe, A. 1983. *Developmental Psycholinguistics: Three Ways of Looking at a Child's Narrative*. New York: Plenum.
- Ravid, D. 2004. Emergence of linguistic complexity in later language development: Evidence from expository text construction. In *Perspectives on Language and Language Development*, D. Ravid & H. Bat-Zeev Shyldkrot (eds), 337–356. Dordrecht: Kluwer.
- Ravid, D. 2006. Semantic-pragmatic development in textual contexts during the school years: The Noun Scale. *Journal of Child Language* 33 (4): 791–821.
- Ravid, D. & Berman, R. A. 2009. Developing linguistic register across text types: The case of Modern Hebrew. *Pragmatics and Cognition* 17: 108–145.
- Ravid, D. & Berman, R. A. In press. Developing noun phrase complexity across adolescence: A text-embedded analysis. *First Language*.
- Ravid, D. & Cahana-Amitay, D. 2005. Verbal and nominal expressions in narrating conflict situations in Hebrew. *Journal of Pragmatics* 37 (2): 157–184.
- Ravid, D. & Katzenberger, I. 1999. Global text evaluation. In *Developing Literacy across Genres, Modalities, and Languages. Working Papers, Volume I*, R. Aisenman (ed.), 220–224. Tel Aviv University: International Literacy Project.
- Reilly, J. S. 1992. How to tell a good story: The intersect of language and affect in children's narratives. *Journal of Narrative and Life History* 2: 355–377.
- Reilly, J. S., Jisa, H., Baruch, E. & Berman, R. A. 2002. Lexical modals in expression of propositional attitudes. *Written Language and Literacy* 5: 183–218.
- Reilly, J. S., Zamora, A. & McGivern, R. F. 2005. Acquiring perspective in English: Use of pronouns, modals, and passives in two genres. *Journal of Pragmatics* 37: 185–208.
- Rumelhart, D. 1975. Notes on a schema for stories. In *Representation and Understanding: Studies in Cognitive Science*, D. G. Bobrow & A. Collins (eds), 211–236. New York: Academic Press.
- Scott, C. M. 1988. Spoken and written syntax. In *Later Language Development: Ages Nine Through Nineteen*, M. Nippold (ed.), 49–95. Austin, Tx: PRO-ED.
- Scinto, L. F. M. 1986. *Written Language and Psychological Development*. New York: Academic Press.
- Slobin, D. I. 1996. From 'thought and language' to 'thinking for speaking'. In *Rethinking Linguistic Relativity*, J. J. Gumperz & S. C. Levinson (eds), 70–96. Cambridge: Cambridge University Press.
- Slobin, D. I. 2003. Language and thought online: Cognitive consequences of linguistic relativity. In *Language in Mind: Advances in the Investigation of Language and Thought*, D. Gentner & S. Goldin-Meadow (eds), 157–191. Cambridge, MA: MIT Press.
- Strömquist, S., Johnansson, V., Kriz, S., Ragnarsdóttir, H., Aisenman, R. & Ravid, D. 2002. Toward a cross-linguistic comparison of lexical quanta in speech and writing. *Written Language and Literacy* 5: 45–69.
- Tolchinsky, L. & Rosado, E. 2005. The effect of literacy, text type, and modality on the use of grammatical means for agency alternation in Spanish. *Journal of Pragmatics* 37: 209–238.
- Tolchinsky, L., Johansson, V. & Zamora, A. 2002. Text openings and closings: Textual autonomy and differentiation. *Written Language and Literacy* 5: 219–254.
- Turner, M. 1996. *The Literary Mind*. Oxford: Oxford University Press.
- Verhoeven, L., Aparici, M., Cahana-Amitay, D., van Hell, J., Kriz, S. & Vigiúé-Simon, A. 2002. Clause packaging in writing and speech: A cross-linguistic developmental analysis. *Written Language and Literacy* 5: 135–162.

Reference points and dominions in narratives

A discourse level exploration of the reference point model of anaphora

Sarah van Vliet

1. Introduction

Sentence and discourse anaphora are traditionally analysed as pertaining to two different domains of linguistic analysis, i.e. syntax and pragmatics respectively.¹ Within Cognitive Linguistics it is assumed that basically the same principles apply both within and across the sentence level. On this view, the unacceptability of full nominals in examples such as (1) and (2) – under a coreferential reading – should be accounted for by the same kind of constraints:

- (1) #*He* put the money in *Steve's* pocket.
- (2) #*He* got up. *Steve* walked to the door.

The account of discourse anaphora presented here is largely based on the *reference point model of anaphora* (Van Hoek 1995, 1997). Van Hoek presents a conceptual-semantic analysis of sentential anaphora, based on the salience signalled by nominal expressions. Nominal elements are defined as *reference points* in semantic *dominions*. Reference points are characterized as local topics, whereas dominions reflect the local context, and are defined as mental spaces (Fauconnier 1994). The acceptability of coreferential full nominals or pronouns depends on referent salience within the local context. Although Van Hoek's model is aimed to replace the structural c-command analysis of *sentential* anaphora, she claims that it can be applied to *discourse* level anaphora as well. In both examples (1) and (2), then, the full nominals are anomalous because they give the wrong signal about the referent's salience within the local context, thus preventing a coreferential reading.

This paper presents a discourse level exploration of the sentential reference point model of anaphora, as applied to the alternating use of proper nouns and pronouns in narratives. The paper is based on examples from English fictional narrative, as well as an experimentally elicited Dutch narrative. It aims to demonstrate that, given a sufficiently

1. I wish to thank Joost Schilperoord, Fons Maes, Leo Noordman, Frederike van der Leek and an anonymous reviewer for helpful comments on earlier versions of this paper.

detailed account of discourse context, an accessibility-based characterization of the nominal category (as in Van Hoek 1997) accounts for the form of reference maintenance in narratives. To this end, I will describe the influence of a number of context factors – such as episode structure, referential distance, point of view and competing referents – on the use of proper nouns and pronouns.

2. The semantics of proper nouns and pronouns

In extending the reference point model of anaphora to the level of (narrative) discourse, I adopt Van Hoek's (1997) semantic characterization of the nominal category, which draws on the notions *accessibility* (Ariel 1988) and *subjectivity* (Langacker 1990). This section presents a brief overview.

During the production and comprehension of discourse, attention flows from one focus to the next and concepts continually move into and out of the immediate focus of consciousness (Chafe 1987, 1994; Langacker 2001). Referential expression types such as (in)definite nominals, proper nouns and pronouns code information pertaining to the current *mental accessibility* of a discourse entity (Chafe 1987; Ariel 1990 inter alia). The nominal categories *definite full nominal* and *pronoun* reflect different degrees of referent salience or accessibility within the immediate context, as represented in the minds of the discourse participants.² Full nominal expressions such as proper nouns are so-called *low accessibility markers*, i.e. they indicate that the intended referent has a low degree of accessibility within the context. Pronouns, on the other hand, are *high accessibility markers*, in that they indicate that the referent is currently highly accessible.³

Van Hoek (1997) distinguishes another aspect of nominal semantics – one that is analogous to the notion of accessibility: full nominals and pronouns differ with respect to the degree of *subjectivity* or *objectivity* they impose on the referent (cf. Langacker 1990). In Van Hoek's words: "Reference via name implies greater distance between the conceived referent and the participants, and a correspondingly more objective conception of the referent. A pronoun portrays the referent as conceptually closer to the discourse participants, and correspondingly as more subjectively construed" (Van Hoek 1997: 219). A subjectively construed element such as a pronoun is a salient part of the assumed conceptual background through which the participants make mental contact with the rest of the conception. In sum, the nominal semantics of proper nouns and pronouns subsumes both accessibility and subjectivity, as "essentially two sides of the same coin" (Van Hoek 1997: 43).

In the present approach to referential form it is assumed that notions such as *accessibility* form an integral part of conceptual-semantic structure. With respect to discourse

2. This characterization includes embedded discourse participants such as story characters (cf. Section 5.6).

3. Although Van Hoek (1997) adopts Ariel's (1988) notion of *accessibility* in characterizing nominal categories, in her own analyses she mostly uses the equivalent term *salience* (and sometimes, *prominence*). When referring to Van Hoek's theory, I will therefore also use the term *salience*.

reference, the proposed accessibility-based characterization of nominal categories is assumed to exert its influence on referential form through *attention framing* (Langacker 2001). As such, the use of proper nouns versus pronouns involves the construal of salience throughout discourse and serves a communicative function.

3. The reference point model of anaphora

This section presents an overview of Van Hoek's (1995, 1997) conceptual-semantic model of anaphora constraints within the sentence. The *reference point model of anaphora* draws on a number of independently developed theories: *Accessibility theory* (Ariel 1990 *inter alia*), *Cognitive Grammar* (Langacker 1987, 1991, henceforth CG), and *Mental Space theory* (Fauconnier 1994).

Within this model, the felicitous use of a coreferential full nominal or pronoun depends on the salience of its referent within the immediate context. In order to plausibly relate (pro)nominal anaphora to the relevant characterization of context, Van Hoek develops a model of semantic relations between nominals, in terms of *reference points* and *dominions*:

Reference points are elements which are prominent within the discourse and so serve to set up the contexts within which the conceptualizer makes mental contact with other entities. The dominion of a reference point consists of the elements that are conceptually located relative to the reference point, whose construal is shaped by their association with the reference point. (Van Hoek 1995: 313)

Reference points, then, function as local topics within a semantic domain, or dominion. A dominion is defined as a *mental space* (Fauconnier 1994), consisting of the (conceptual, non-linguistic) context within which a particular referent is the most salient element. A pronoun can only be used if there is a salient antecedent that can function as reference point for the interpretation of that pronoun. Full nominals are used if their relevant context of interpretation falls outside the dominion of a corresponding reference point. Using a full nominal that necessarily falls *within* the dominion of a corresponding reference point sends the wrong signal concerning the referent's retrievability within the context, preventing an intended coreference reading.

Sentential *reference point / dominion organization* depends on the salience of a nominal and (the strength of) its conceptual-semantic connections with coreferential nominals. I will now describe the factors that shape this organization, as distinguished by Van Hoek: (i) conceptual connectivity; (ii) linear order; and (iii) point of view. In the subsequent Sections 4 and 5 these same factors will be applied to the level of narrative discourse.

3.1 Conceptual connectivity

Conceptual connectivity pertains to the strength of relations between nominal elements, be it within a single clause, a sentence or an entire discourse. This conceptual connectivity

partly determines the possibility for a referent to 'escape' the dominion of a coreferential nominal, and to be coded by a full nominal.

Conceptual connections within the clause are reflected in the grammatical relations between verb, complements, and modifiers. Note that CG views such syntactic relations as surface manifestations of underlying conceptual-semantic structure (cf. van Hoek 1997: 10, 66). Van Hoek uses the notion of *complement chain*, which is roughly analogous to the grammatical relations hierarchy: subject > direct object > indirect object > oblique (Keenan and Comrie 1977 *inter alia*). This chain reflects the hierarchy of prominence within the central clause, in which the subject is more prominent than the object, and so forth. Outside the complement chain, the nominals within various types of modifiers represent the least prominent entities with respect to other entities in the clause. Constraints on coreference within the clause largely follow from the prominence asymmetries reflected in these patterns.

To give an example, the *subject* is the most prominent element within the complement chain, and therefore functions as the reference point for the central process described by the clause. This explains the unacceptability of sentence (4) under a coreferential reading:

(3) *John* likes *his* mother

(4) **He* likes *John's* mother

(Van Hoek 1997:65)

Van Hoek explains that, the subject being the most salient entity in the sentence, *he* in (4) is the *main reference point* within the complement chain. Any corresponding nominal within that chain must be highly salient as well, because it is conceptualized *relative* to the subject reference point. Consequently, the use of a proper noun, which signals low accessibility, prevents a coreferential reading.

Van Hoek argues that clausal conceptual connectivity is not just a notational variant of syntactic c-command relations. This can be illustrated by the following examples:

(5) **Mary* hit *him* just before *John* got up

(6) *Mary* hit *him* before *John* had a chance to get up

(Brugman and Lakoff 1987, cited in Van Hoek 1997:92)

The first example is ungrammatical under a coreferential reading. In the second sentence the conceptual break improves coreference possibilities. That is, the *irrealis* subclause allows the referent *John* to fall outside the dominion of the coreferential object nominal. Examples like these demonstrate the fundamentally *semantic* nature of the anaphora constraints, which allows for the possibility that sentential and discourse anaphora be accounted for in the same way.

Moreover, conceptual connectivity between nominal elements represents a *continuum* ranging from the strong connectivity between complements of the verb, to weaker connectivity found with clausal modifiers, to weakest connectivity, reflected in discourse unit boundaries. This weakest type of conceptual connectivity will be further discussed in Sections 4 and 5.

3.2 Linear order

Another factor in reference point / dominion organization is the *linear order* of corresponding nominals. Van Hoek posits that, in general, “a nominal tends to be construed as a reference point in relation to elements which follow it in the linear string, and is less likely to be construed as a reference point in relation to elements which precede it” (Van Hoek 1997:80). The importance of this factor, however, is “in inverse proportion to the strength of the connection between two entities” (Van Hoek 1997:227). The influence of linear order can be seen by comparing sentences (7) and (8); the clause-initial position of the modifier in (7) improves the acceptability of the full nominal.

(7) In *Carter’s* home town they still consider *him* a genius.

(8) They still consider *him* a genius in *Carter’s* home town. (Van Hoek 1997:99)

Note that, as these examples illustrate, grammaticality judgments are considered a matter of degree rather than a binary distinction.

3.3 Point of view

Van Hoek defines dominions as *mental spaces*, “domains that we set up as we talk or listen, and that we structure with elements, roles, strategies and relations” (Fauconnier 1994: 1). The construal of a new point of view (POV) opens up a new mental space. The use of (pro)nominal anaphora may then depend on the extent to which the material is construed from the POV of a referent. If an utterance is interpreted as being in a cognizing or viewing relation with a (previously mentioned) referent, a corresponding entity within that utterance falls within that referent’s dominion, triggering pronominal reference. Since viewing or cognizing relations are often implicit, it is sometimes possible to impose different mental space configurations onto a given utterance, i.e. to construe it as either conceived by the speaker or by a referent.

(9) That *he* was blond worried *John*

(10) That *John* was blond worried *him*

(11) *That *John* was always unhappy worried *him* (Van Hoek 1997:209)

In (9) the subclause is interpreted as representing *John’s* conception. The subject of the subclause therefore falls within the semantic dominion of the referent *John*, and is pronominalized. As can be observed in sentence (10), a subclause such as the one in (9) can also be construed objectively; that is, since features such as hair colour are perceptible to everyone, the subclause – although part of *John’s* conception – can be construed from a POV other than *John’s*. In (10) therefore, the subclause *That John was blond* falls outside the dominion of the main clause object, licensing the use of a proper noun in the subclause. Sentence (11) demonstrates that a difference in POV may lead to differences in acceptability in structurally similar sentences: predications about perpetual inner states of mind cannot be construed from an objective POV, but must reflect the referent’s own

conceptualization. Since the referent *John* in (11) cannot but be construed as conceptualizer, it cannot escape the dominion of the coreferential object, and should therefore be coded by a pronoun.

To summarize, conceptual connectivity, linear order and POV are the conceptual-semantic factors affecting sentential reference point / dominion organization, and thereby the sentential anaphora constraints.

4. A discourse level reference point model of anaphora

Van Hoek (1997, Chapter 5) describes a number of empirical studies involving discourse level referential patterns, which she suggests can be related to the reference point model. The remainder of this paper is concerned with extending the model to the level of narrative discourse. It aims to demonstrate that discourse level referential patterns can indeed be described using the same factors as in the sentential reference point model, namely conceptual connectivity, linear order and point of view. Section 5 describes how these factors operate in narrative discourse. This section addresses the reasons for the discourse level extension (§4.1), and presents a sketch of what narrative reference point / dominion organization might look like (§4.2).

4.1 Reasons for extending the reference point model to the level of discourse

The separate treatment of sentence and discourse anaphora is partly due to the different research methods used. Whereas discourse level research focuses on the distribution of referential form (e.g. through corpus analysis), sentence level analysis is mainly concerned with grammaticality judgments. This sentence/discourse distinction is supported by the observation that sentential anaphora tend to evoke strong grammaticality judgments (cf. Lasnik 1989). However, this does not necessarily imply that there is a principled distinction between sentence and discourse anaphora, or that they obey fundamentally different constraints. As Van Hoek argues, the difference in acceptability judgments may be one of degree rather than kind, and may reflect “the difference between overtly coded versus implicit semantic interconnections” (Van Hoek 1997: 13). According to Van Hoek, “[t]he strongest connectivity, involving head/complement relations within clauses, gives rise to the most unequivocal judgments because there is little or no flexibility in the construal of the relationships between the nominals. In multisentential discourse, there is more room for alternate construals” (Van Hoek 1997: 13). Sentential patterns of usage are more entrenched and less variable than discourse patterns, but constraints on (pro)nominal anaphora resulting from the conceptual organization reflected in these sentential patterns need not be *qualitatively* different from constraints on discourse anaphora.

A unified treatment is supported by the fact that the same morphological forms are involved; topic referents in narrative discourse tend to display a proper noun/pronoun alternation pattern. Van Hoek convincingly argues that constraints on sentential (pro)nominal anaphora need not be stipulated separately in terms of structural rules, but

fall out naturally from (i) the notion of accessibility inherent in nominal semantics, and (ii) the clausal prominence relations described in CG. Given this view of nominal semantics – combined with the conceptual-semantic rather than structural characterization of nominal contexts – the same type of constraints may apply to the use of anaphoric proper nouns and pronouns over longer stretches of discourse. The reference point model offers the theoretical constructs necessary for an accurate characterization of context factors affecting referent accessibility and referential form in narrative discourse. A discourse level extension is further supported by the nature of referential dominions as mental spaces, which allows for embedding and thereby the inclusion of discourse dominions.

4.2 A sketch of reference point / dominion organization at the level of narrative discourse

On the view of anaphora presented here, the factors involved in discourse level reference point / dominion organization are essentially the same as within sentences – they just involve global as well as local salience, and conceptual connections exert their influence over longer stretches of text. The dominion of a sentential subject, for instance, may extend beyond the immediate clausal context, provided that there is no break in conceptual connectivity or other factor which closes off the current dominion, and triggers a repeated proper noun.

Discourse dominions consist of the conceptual (not necessarily linguistic) structures that are (to be) interpreted relative to a local or global discourse topic, in narrative typically a character. A discourse dominion can perhaps best be thought of as the conceptual representation of a (fictional) situation, in which a single element is most salient or accessible, and functions as a conceptual reference point for that situation. Within the current reference point approach, the extent of a referential dominion, and thereby the extent of continued pronominalisation, is delimited by the factors' linear order, conceptual connectivity and point of view, be it at the clause or discourse level.

An account of discourse level anaphora must take into account the time course and, consequently, the flow of attention throughout discourse. However, referent accessibility cannot be directly equated with concept activation: referential form is tailored to assumed accessibility for the hearer/ reader. For the narrator, of course, a referent is always highly accessible by the time it is ready to be verbalized. Within the current accessibility-based approach, the use of proper nouns versus pronouns throughout discourse is part of what Langacker calls "attentional framing" (2001: 154ff.). According to Langacker, "[t]his dimension of linguistic organization involves the presentation of information in coherent "packets" of digestible size" (2001: 154). Its semantic contribution lies in imposing a "window of attention" on the conceptual content. Through attention framing, then, the choice between (repeated) proper noun and pronoun contributes to the construal of referent salience and the packaging of information in discourse.

For narratives, an important determinant of reference point status is the inherent salience of discourse topics such as characters. In addition, it is determined by the conceptual-semantic connectivity between a nominal referent and a corresponding entity within the narrative context. An account of the conceptual structure of narratives may therefore

shed light on proper noun/ pronoun patterns. Underlying this approach is the assumption that the narrative is a cognitively relevant category: *stories* play a central role in human culture, cognition and language (Turner 1996). Another important assumption is that the representation of narratives (in terms of dominions) is primarily a conceptual rather than linguistic construct (cf. Fauconnier 1994).

The organization of reference points and dominions in narratives can be illustrated by the following example:

- (12) *Francis* got home late from town, and *Julia* got the sitter while *he* dressed, and then hurried *him* out of the house. The party was small and pleasant, and *Francis* settled down to enjoy *himself*. (J.C.CH: 390)⁴

In the first sentence, the protagonist subject (*Francis*) retains its status as primary reference point, and remains pronominalized, even though there are two intervening referents. In the second sentence the break in conceptual connectivity (in terms of location, time, activity and implied other characters) closes off the referential dominion. In addition, the intervention of the local topic *Julia* (in the previous sentence) diminishes the accessibility of the primary reference point *Francis*. Within the new context the protagonist is no longer as accessible. As a consequence, a new referential dominion is created and a low accessibility full nominal is used to refer to this referent within its new context.

Note that within the representation of a single clause, several referents may serve as reference point, as can be observed from the use of both null subject and pronoun in the clause *and hurried him out of the house*. Most stretches of discourse involve various different topics, and several simultaneously active referential dominions can be embedded in one another.

Lastly, it must be noted that reference point status is not the only factor involved in referential form. In some cases, such as alternative descriptions, indirect anaphora, evolutive referents (Maes 2001), certain uses of demonstratives (Maes and Noordman 1995) or metonymic descriptions, the referential expression itself adds information about the referent, and its form is determined by other factors than accessibility (alone). I propose therefore that the scope of reference point / dominion organization is *reference maintenance* throughout discourse.

5. Reference point / dominion organization in narratives

This section systematically relates the factors distinguished in the reference point model to referential patterns in English fictional narrative. The relevant factors – salience, conceptual connectivity, linear order, and point of view – are applied to a narrative discourse context rather than a clausal context. It is shown that several grammatical and discourse factors that can be captured in terms of the reference point model indeed affect the alternation of proper nouns and pronouns in reference to narrative characters.

4. Examples taken from popular and literary fiction are labelled by abbreviations. The sources are listed at the end of the paper.

The set-up of the section is as follows: Section 5.1 discusses the inherent *salience* of narrative protagonists; in Section 5.2 referential distance is described as a discourse level counterpart of *linear order*; Section 5.3 describes how intervening referents may close off a referential dominion; the next two Sections, 5.4 and 5.5, are concerned with *conceptual connectivity*, as reflected in clause structure and narrative structure respectively; Section 5.6 describes the influence of *point of view* in narratives; Section 5.7, lastly, illustrates how the various factors interact at any given point during the flow of discourse.

5.1 Protagonists and the proper noun/pronoun pattern

At the level of discourse, we may distinguish global and local accessibility. Story characters are natural candidates for reference point status within the global narrative representation: they are the prime ‘anchors’ to which the story is connected. Particularly protagonists may retain their reference point status throughout longer stretches of text. Karmiloff-Smith (1981) for example found that whereas narrators often use pronouns to refer to the main character, secondary characters are often coded by full noun phrases, even when they have been mentioned more recently.⁵

Since a global discourse topic may remain accessible throughout the discourse, several referential dominions may be included in the context of a *main* discourse reference point. For example in a story or episode about a single protagonist, even parts that do not directly involve the protagonist are in some sense interpreted relative to this character, and often also from his/her POV. The following excerpt illustrates how global reference point status triggers the continued use of pronouns. The relevant chapter opens with a lengthy paragraph describing the U.S. attorney Roy Foltrigg. The second paragraph starts a new episode and runs as follows:

- (13) As *he* [Roy Foltrigg] entered the Federal Building on Main Street in Memphis, a few minutes after midnight, *he* had an escort of sorts with Wally and Fink and agents Trumann and Scherff, but there were no anxious reporters. In fact, not a soul waited for *him* until *he* entered the offices of the FBI where Jason McThune sipped stale coffee with two other weary agents. So much for grand entrances. (JG.TC: 73)

The example shows how reference to a topic character is maintained through pronouns throughout an entire paragraph, despite factors that might otherwise have triggered a repeated name, such as the onset of a new episode, various intervening characters and the switch to syntactic object position (in the second sentence). Also note the fragment *so much for grand entrances*, which seems to be a representation of Foltrigg’s own thoughts, reinforcing the reference point status of this character.

In spite of their natural reference point status, central characters in narratives are not invariably pronominalized, but display an alternation pattern of pronouns and repeated proper nouns; the assignment of reference point status at a specific point within

5. Van Hoek (1997) suggests that the asymmetry between primary and secondary characters can be viewed as analogous to the subject/object asymmetry in clausal relations.

the discourse depends not only on the global salience of the character, but also on local context factors. These factors will be addressed in the following sections.

5.2 Referential distance

One of the factors determining reference point / dominion organization is *linear order*; a referent typically functions as a reference point for corresponding nominals which follow it in the linear string. As I understand it, this is a natural consequence of limited attention: within sentences, a pronoun followed by a corresponding proper noun requires more cognitive effort than a noun followed by a pronoun. If we interpret the *sentential linear order* factor as a reflection of such attention constraints, its *discourse* counterpart is *linear distance*.

It is well known that referential distance affects referential form. Givon (1983) for example found that an increase in the number of clauses between mentions results in a higher proportion of full nominals. The crucial factor in distance-based repeated full nominals seems to be *delay of concept activation* (cf. a.o. Deane 1991): the activation of a mental entity that is not focused on for some time gradually diminishes in the working memory of the discourse participants. Consequently, the entity's referential dominion is closed off because it is no longer salient, and resumed reference triggers the use of a full nominal. The effect of referential distance may be interpreted in two ways: in terms of the intervening processing time between consecutive references, and in terms of intervening information, such as other characters, concepts, and conceptualized events (cf. 5.3 below).

Referential distance may also work the other way around, and *prevent* the use of a full nominal. In the following example, the use of two or three names in a row is infelicitous and might even suggest that somehow different 'Matildas' are involved:

- (14) *Matilda* stood up and \emptyset ?*Matilda* began to say the two-times table. When *she*?*Matilda* got to twice twelve is twenty-four *she* didn't stop. *She* went right on with twice thirteen is twenty-six, twice fourteen is twenty-eight, twice fifteen is thirty, twice sixteen is...

(RD.M:64)

The use of two coreferential full nominals directly following each other in closely connected sentences sends the wrong signal about the referent's current high accessibility. Gordon et al. (1993) also found that this tends to slow down reading times (the so-called *repeated name penalty*).

It seems intuitively plausible that the influence of referential distance on referential form is part of attention framing. I propose that other, more clearly semantic factors (such as episodic structure) are also related to this windowing of attention. In spontaneous discourse, the various relevant factors are assumed to exert their influence on referential form 'in tandem', as part of attention framing.

5.3 Competing referents

One of the factors that may close off a dominion is the intervention of another referent. In cases of same gender/number referents, repeated proper nouns are used in order to avoid ambiguity for the hearer/reader. However, the very presence of another referent – whether or not it differs in number/gender – tends, especially if it functions as a (local) topic, to diminish the accessibility of the intended referent and close off the referential dominion. The influence of referent competition can be illustrated by the following excerpt:

- (15) *Strauss* hits a fast serve, and once again it's a body line, aimed straight for the shoulder. *Perowne* manages to push his racket through the ball, and the volley goes more or less as he hoped, and now he's in position, on the 'T'. *Strauss* flicks the ball out of the corner, and it comes back along the same sidewall. *Perowne* goes forward and volleys again. (IME.S:108-9)

Another factor that influences the extent to which competing referents trigger repeated proper nouns is their salience within the clause, as reflected in their syntactic position. This will be addressed in Section 5.4 below.

5.4 Salience within the clause

As noted by Van Hoek, the conceptual-semantic prominence hierarchy reflected in grammatical relations (cf. 3.1 above) affects coreference possibilities within the clause. It also has consequences for the extent of a reference point's dominion, and consequently referential form, when the antecedent is not contained within the same clause.

Throughout discourse, subjects are more likely to be pronominalized than other referents (Chafe 1976; Kuno 1987 *inter alia*). In centering theory for example (Grosz et al. 1995) it is observed that for consecutive clauses, subjecthood is an important factor affecting the continuation of topic status and pronominal reference. Chafe (1994) claims that subjects typically represent starting points for the clause. Such starting points are usually accessible referents, which may explain the close association between subjecthood and pronominalization.

The tendency for repeated use of proper nouns is expected to increase further down the grammatical relations hierarchy. That is, obliques are more likely to be full nominals than indirect objects, which in turn are more likely to be full nominals than direct objects. This factor can be illustrated by the example below; the shift to oblique complement is accompanied by a resumption of the proper name:

- (16) As *he* straightens up, it occurs to *Perowne* that what *he* really wants is to go home and lie down in the bedroom and think it through, the dispute in University Street, and decide how *he* should have handled it, and what it was *he* got wrong. (IME.S:102)

The prominence asymmetry reflected in the complement chain may also affect the extent to which a *competing referent* triggers the repeated use of proper nouns for an established topic. An intervening *object* is less likely to diminish the global topic's reference point status than an intervening *subject*. An intervening *indirect object* or *oblique* referent is even

less likely to affect the reference point status of the topic. In the following example, the intervention of a local topic subject (another character, *Baxter*, underlined) triggers the repetition of a proper noun when reference to the main character is resumed:

- (17) Above all, there swells in *him* [Henry] a peculiarly modern emotion – the motorist’s rectitude, spot-welding a passion for justice to the thrill of hatred, in the service of which various worn phrases tumble through *his* thoughts, revitalised, cleansed of cliché: just pulled out, no signal, stupid bastard, didn’t even look, what’s his mirror for, fucking *bastard* [italics in original]. The only person in the world he hates [Baxter] is sitting in the car behind, and *Henry* is going to have to talk to him, confront him, exchange insurance details with him – all this when *he* could be playing squash. (IME.S: 82)

This paragraph represents a protagonist’s inner monologue, which reinforces its reference point status. A competing subject referent, however, triggers a repeated proper noun for subsequent reference to the protagonist. Note that within the underlined phrase constituting the intervening subject, the discourse topic *itself* functions as a reference point to identify the intervening character. Within the larger context of the entire clause, however, the intervening subject (rather than the relative clause subject within it) is the primary reference point.⁶ Return to the protagonist referent in the next clause is therefore accompanied by resumption of the proper noun *Henry*. Subsequently, the intervening character *Baxter* is again mentioned three times, in object and oblique position. This, however, does not trigger a repeated proper noun in the last clause, since the subject referent continues to be *Henry*.

The next section addresses the influence of another type of conceptual connectivity on reference point / dominion organization, i.e. narrative structure.

5.5 Narrative structure

5.5.1 Episode structure

For narratives, the main discourse level correlate of conceptual connectivity is episode structure. According to Van Dijk (1982), the episode plays an important role in storage and retrieval of discourse information, in marking the difference between more and less important information, and in the overall organization of a coherent discourse representation. Schilperoord’s (1996) study of written text production provides empirical evidence that units such as paragraphs reflect cognitive planning units on the part of the discourse producer. For narrative comprehension, there seems to be ample experimental evidence for the psychological reality of the representation of described events, and for the importance of conceptual connectivity between those events. Readers routinely keep track of (the continuity of) the events described in the text, rather than only the sum of propositions contained in the text. According to Zwaan et al. (1995), successive events are *indexed* along the dimensions of *time*, *space*, *causation*, *motivation* and *protagonist*. The (dis)continuity of events in terms of these dimensions affects the integration of successive

6. Cf. also Van Hoek’s discussion on reference points at various levels of conceptual organization (1997: 67ff.)

clauses within the overall narrative representation; each of these dimensions has been demonstrated to affect for example reading times and memory retrieval (cf. Zwaan and Radvansky 1998 and references therein).

These empirical observations seem to reflect the notion of conceptual connectivity in narrative discourse. They also support its characterization as a continuum; consecutive clauses represent either a break or a continuation on any of these dimensions. In addition, the breaks themselves may differ in strength (e.g. a flashback to a description of a previous century represents a greater conceptual break than a transition to 'the next day'). In view of its composite nature, then, episode structure is not necessarily an all-or-nothing affair.⁷ Further, the conceptual rather than linguistic status of represented events is consistent with the current approach to narrative representation.

I adopt the following working characterization of episode structure: Episode structure is the reflection of conceptual connectivity in narratives. It is a continuum comprising (at least) the parameters *character*, *cause*, *motivation*, *location* and *time*.

- Episode breaks occur when there is a (weak or strong) change in one or more of these parameters that define the conceptual connectivity throughout the narrative. Consequently, episode boundaries represent conceptual shifts of variable degree and kind.
- There are various text signals for the demarcation of episodes, such as pauses, paragraph indentations, and place and time markers, amongst others (Van Dijk 1982: 181).

The nature of episode structure can be illustrated by some examples. The episode break in (18) represents a strong conceptual break; there is a shift in time, location, characters, and cause: from a description of a safe and pleasant childhood to an unexpected terrible event. The episode transitions in example (19), on the other hand, involve only a temporal break and a continuation in other respects.

- (18) Until he was four years old, James Henry Trotter had a happy life. He lived peacefully with his mother and father in a beautiful house beside the sea. There were always plenty of other children for him to play with, and there was a sandy beach for him to run about on, and the ocean to paddle in. It was the perfect life for a small boy.

Then, one day, James's mother and father went to London to do some shopping, and there a terrible thing happened. Both of them suddenly got eaten up (in full daylight, mind you, and on a crowded street) by an enormous angry rhinoceros which had escaped from the London zoo. (RD.JGP:7)

- (19) The two women and the small boy stood absolutely still on the grass underneath the tree, gazing up at this extraordinary fruit. James's little face was glowing with excitement, his eyes were as big and bright as two stars. He could see the peach swelling larger and larger as clearly as if it were a balloon being blown up.

In half a minute, it was the size of a melon!

In another half-minute, it was twice as big again!

(RD.JGP:21)

7. However, in written discourse a decision as to episode structure is often forced because of orthographic structure.

5.5.2 *Episodes and reference point / dominion organization*

The extent of referential dominions is partly determined by the conceptual connectivity reflected in episodic structure; an episode transition may close off the current referential dominion, triggering a repeated full nominal, even if the same character appears in both episodes. The repetition of full nominals after episode boundaries has been demonstrated in various empirical studies (a.o. Clancy 1980; Chafe 1980; Marslen-Wilson et al. 1982; Fox 1987; Tomlin 1987; Vonk et al. 1992). Within Van Hoek's reference point model, this tendency can be explained as follows: first, the cognitive effort involved in the conceptualization of a new episode causes a disruption in the flow of attention, and leads to a lower degree of referent activation and thereby reference point status. A related explanation involves the notion that the extent of a reference point's dominion depends on conceptual connectivity between coreferential nominals; shifts along the dimensions of time, space etc. disrupt conceptual connections between a nominal and its antecedent contained within the previous episode. In other words, these conceptual connections are weaker if the referents are embedded within different settings.

The tendency to repeat proper nouns at the onset of new episodes can be illustrated by the following example:

- (20) Nick hastily finished his own drink, and said, 'Thanks. Or maybe this time I'll have a shot of rum in it.'

After half an hour more *Nick* had slid into a kind of excited trance brought on by *his* new friend's presence [...]
(AH.LB: 33)

Note that a pronoun would be considered perfectly acceptable in this sequence. In fact, there is great variability in the tendency to repeat proper nouns (rather than maintain pronouns) after episode breaks. This can be explained by the following two factors:

First, reference point / dominion organization is not about conceptual connectivity alone; it is about the *interaction* of (global) referent accessibility and conceptual connectivity. There is fluctuation of referent accessibility throughout the text, as well as variation in the salience of different referents relative to each other. That is why inherently topical entities such as protagonists are often pronominalized throughout, even after episode boundaries. Other factors that may prevent repetition of proper nouns are e.g. the need to avoid two proper nouns immediately following each other, or subjective construal of a topic referent.

Further, as pointed out above, conceptual connectivity between episodes is variable, and the tendency to repeat proper nouns is influenced by the extent to which consecutive episodes are viewed as being closely related. This can be illustrated by comparing example (20) to example (21) below:

- (21) *Nick* was glad *he* wasn't going to Nat's wedding, and yet *his* absence, to anyone who noticed, might seem like an admission of guilt, or unworthiness. *He* saw a clear sequence, like a loop of film, of *his* friends not noticing *his* absence, jumping up from gilt chairs to join in the swirl of a ball. On analysis *he* thought it was probably a scene from a Merchant Ivory film. The doorbell trilled and *Nick*/? *he* saw a van in the street where the Rolls had been. *He* went out and there was a skinny boy in a baseball cap pacing about, and some very loud music.

(AH.LB: 486-7)

In this excerpt from the same novel, involving the same protagonist, the episode shift is also accompanied by a repeated proper noun. Here, however, a pronoun at the episode onset would be less felicitous than in (20), since the break in conceptual connectivity is much stronger: the first episode describes the character's fantasy, of an envisaged situation remote from the current story setting, At the episode boundary an external cause (the bell) shifts the story back to the main setting, the reality of the story. The break in conceptual connectivity therefore involves not only time, but also location. In addition, there is a shift from the character's inner world to external reality.⁸

Recall that acceptability judgments such as the ones described above are considered a matter of degree rather than kind, even at the clause level. We may conclude that the variability in episode-related proper noun repetition does not imply that discourse anaphora constraints are fundamentally different from sentence constraints.

These variable possibilities in turn enable narrators to communicatively *exploit* reference point / dominion organization for structuring the discourse. That is, the narrator is able to influence the construal of conceptual connectivity between episodes by the choice of referential form: establishing a new referential dominion by means of repeated full reference reinforces the conceptual break between consecutive episodes; the continuation of a referential dominion through pronominalisation, on the other hand, enhances the conceptual connectivity between the episodes.

5.5.3 *Background propositions and plot-advancing propositions*

An additional aspect of narrative structure relevant to referential patterns is the general distinction between story *background* and plot *advancement* (a.o. Werth 1999). Background propositions constitute the conceptual background relative to which the story makes sense; plot-advancing propositions specify the actions and processes taking place within the story. The break in conceptual connectivity involved in going from plot to background may close off the current referential dominion and trigger a repeated full nominal such as a proper noun, as in the following example:

- (22) *Matilda* longed for *her* parents to be good and loving and understanding and honourable and intelligent. The fact that they were none of these things was something *she* had to put up with. It was not easy to do so. But the new game *she* had invented of punishing one or both of them each time they were beastly to *her* made *her* life more or less bearable. Being very small and very young, the only power *Matilda* had over anyone in *her* family was brainpower. For sheer cleverness *she* could run rings around them all. (RD.M: 43)

One of the properties of background propositions is that their predicates denote inherent *states* and *properties* rather than actions or processes. This characteristic also affects referential form; as Bolinger (1979) points out, repeated full nominals may emphasize

8. The variable strength of episode boundaries, and the accompanying difference in the tendency to repeat proper nouns, can be further illustrated by comparing example (20) to example (12) in 4.2 above. The conceptual break in (12) involves more event parameters than the one in (20): a shift involving time (inferred), location, motivation (i.e. the character's goals) and character (i.e. the inferred presence of other characters). Note that in (12), pronoun continuation would be less acceptable than in example (20).

the inherent nature of the referent: “X qua X; X has the quality suggested by the clause in which X occurs” (Bolinger 1979: 291). The following is Bolinger’s example (ibid.):

- (23) You don’t need *sulfur* for drying apricots; *sulfur* ruins the flavor.

In narratives, such ‘X qua X’ sentences are typically used for background propositions. I propose that within the larger discourse context, conceptual connectivity is relevant to such cases as well. If the background proposition is continued by other background propositions, forming part of an elaborate description, repeated references fall within the background dominion and are further pronominalized. This can be illustrated by the following example:⁹

- (24) Aunt Sponge was enormously fat and very short. She had small piggy eyes, a sunken mouth, and one of those white flabby faces that looked exactly as though it had been boiled. She was like a great white soggy overboiled cabbage. Aunt Spiker, on the other hand, was lean and tall and bony, and she wore steel-rimmed spectacles that fixed onto the end of her nose with a clip. She had a screeching voice and long wet narrow lips, and whenever she got angry or excited, little flecks of spit would come shooting out of her mouth as she talked.

(RD.JGP: 11)

Conceptual discontinuities can also occur in plot-advancing propositions, which are part of the central story. Fox (1987) finds that in written narratives repeated full nominals are often used in *development structures*, when a character initiates an *action*, in reaction to foregoing events (cf. Fox 1987 for similar examples):

- (25) But then suddenly, just as *he* [James] was passing underneath the old peach tree that stood in the middle of the garden, *his* foot slipped, and *he* fell flat on his face in the grass. The paper bag burst open as it hit the ground and thousands of tiny green things were scattered in all directions.

James immediately picked *himself* up onto *his* hands and knees, and started searching around for *his* precious treasures.

(RD.JGP: 20-21)

Within the current composite approach to episodic structure, development structures can be said to involve the dimensions of *cause* and *motivation*. They differ from prototypical episode boundaries in that they do not involve temporal and locative shifts, but can be accounted for in the same way, namely in terms of the continuum of conceptual connectivity.

Apart from the narrative factors described above, reference point status is affected by discourse perspective. This will be addressed in the next section.

9. Note that in this example, the use of the proper noun *Aunt Spiker*, although part of a property description and background proposition, also involves disambiguation.

5.6 Subjective and objective contexts

Perspective is an important and pervasive characteristic of all narrative discourse. For current purposes I adopt Sanders' (1994) definition of discourse perspective as *subjective point of view*:

Perspective is the introduction of a subjective point of view that ascribes the claim of validity of the presented information to a particular subject (person) in the discourse. A discourse segment is perspectivized if its relevant context of interpretation is a person-bound, embedded space within the narrator's reality. (Sanders 1994: 37)

In various degrees of directness, the responsibility for content and verbalization of speech and thought can be attributed to sources other than the narrator. A so-called *implicit perspective* represents the character POV in a subtle way and may be achieved by a variety of linguistic means such as modal verbs, verbs of cognition and perception, and referential expressions. This can be illustrated by the following newspaper example from Sanders (1994: 59):

- (26) The police lost track of the car with *the kidnapped girl*. In the woods near Apeldoorn, a policeman discovered a man who had *a girl* with him. The kidnapper had released *her* in a nearby street.

In this example the italicized referent is the central character of the story. Nevertheless, the switch to another character's POV triggers the use of an indefinite nominal, reflecting the newness of the topic within the perspectivized space.

Sanders and Redeker (1996) propose that Langacker's (1990) notion of subjectivity be extended to include character subjectivity. In this way, a referent's information status is assessed not only relative to the discourse participants, but also relative to the embedded character.¹⁰

The relation between this character subjectivity and reference point / dominion organization will be described below.

5.6.1 Subjectivity and the proper noun/pronoun distinction

Given the subjective nature of pronouns, they are often used in reference to characters that function as conceptualizer of the propositional content, as in example (27) below:

- (27) *He* needed a new lawyer, one who would return *his* phone calls and meet *him* for drinks and find some jurors who could be bought. A real lawyer!
He needed a new lawyer, and *he* needed a continuance or a postponement or a delay, hell, anything to slow this thing down so *he* could think. (JG:TC: 27)

This excerpt represents the character's own thoughts, as reflected in the exclamations *a real lawyer* and *hell*, and in the uncertain legal descriptions. A proper name in the last sentence would be anomalous, because it would be incompatible with the referent's status

10. Here I use the more general term 'information status' instead of reference point status, because the indefinite nominal involves (embedded) *identifiability* and *newness*.

as the conceptualizer. Note that the *sentential* (pro)nouns in POV examples (9) through to (11) (Section 3.3 above) are accounted for in the same way; the name in sentence (11) is incompatible with its reference point status as experiencer of the predication. This confirms that the same POV constraint on (pro)nominal anaphora holds both within and across sentences.

In some cases, subjective construal may even override accessibility for the discourse participants; Emmott (2003) observes that a protagonist's perspective may present another, envisaged, character as accessible, and license the use of pronouns, even when a third character has been mentioned more recently. The following is Emmott's example from a popular fiction novel (her example 2, p. 298):

- (28) I went to bed early ... When *Jake* climbed in beside me later, I pretended to be asleep, though I lay awake for hours in the dark. I planned what I would wear. I thought about how I would hold **him** [Adam], learn his body, trace his ribs and his vertebrae...

Whereas *pronouns* impose a subjective construal, *full nominals* such as proper nouns present a more objective construal of the referent (Van Hoek 1997, cf. Section 2 above). Bolinger (1979) notes that repeated full nominals "may involve an extraneous viewpoint whereby the speaker attributes to the referent some expression that is not (or not entirely) the referent's own at the time: the referent looking at himself, some point of general information, or an opinion of the speaker – a sort of concealed quotation" (Bolinger 1979: 308). In the following example, the repeated proper noun presents an objective construal of the referent, as presented from the POV of a secondary character:

- (29) *Jimmy* knew what burning hair smelled like because *he'd* cut off some of *his* own hair with the manicure scissors and set fire to it with *his* mother's cigarette lighter. [...] *His* father had laughed then, but *his* mother hadn't. At least (*his* father said) *Jimmy'd* had the good sense to cut the hair off before torching it. (MA.OC: 18)

As illustrated above, pronouns are associated with subjective construal, and proper nouns with objective construal. In many cases, it may be difficult to disentangle subjectivity from accessibility. Note that this is not problematic for the present schematized characterization of nominal semantics (cf. Section 2 above). The influence of subjectivity specifically can be observed when a narrator uses both a pronoun and a proper noun in juxtaposition, as in example (30) below:

- (30) The three people in the world **he**, Henry Perowne, most loves, and who most love **him**, are about to come home. So what's wrong with him? Nothing, nothing at all. He's fine, everything is fine. (IME.S: 181)

In this excerpt, the pronoun is used, arguably, because the protagonist is portrayed as the experiencer, reinforced by the use of free indirect discourse (witness the question and repetitions). The full nominal used directly after the pronoun may indicate that this character is looking at himself 'from the outside', trying to analyse himself as objectively as possible. The use of both pronoun and proper noun directly following each other can be ascribed to the incorporation of both subjective and objective construal.

5.7 Interaction of factors

Particularly at the level of discourse, reference point / dominion organization reflects the *interaction* of various factors. For example, a narrator may be dealing with conflicting factors: an episode transition occurring just *after* a full nominal sometimes prevents the use of yet another proper noun. In addition, the subjective construal of a referent may prevent the use of a repeated name after an episode transition (as in example (27) above). Furthermore, factors such as accessibility and conceptual connectivity are presumably not represented as binary features, but rather as a continuum. This interaction of factors causes variability in referential patterns. That is, none of these factors can be defined as a direct trigger or rule for the assignment of referential form.

The following example illustrates the interaction of some of the aforementioned factors:

- (31) From Gary and his wife, in addition to the port, *Chip* received a clever vacuum-pump system for preserving leftover wine from oxidation, as if leftover wine were a problem *Chip* had ever had. (J.F.T.C.:82)

In spite of the relatively small referential distance, the second mention of *Chip* is by a full nominal. The repeated name emphasizes the inherent nature of the character, as in Bolinger's (1979) 'X qua X' sentences. In addition, the full nominal reflects the break in conceptual connectivity through the introduction of the *irrealis* context (a *counterfactual* space in the sense of Fauconnier 1994).

Discourse level referential choices involve the *interaction* of various factors affecting referent accessibility within the current context. This is one of the reasons that more detailed distributional analyses of discourse reference are also needed in order to confirm the validity of the various individual factors discussed above.

6. The distribution of proper nouns and pronouns in written Dutch narratives

The literary examples given in this paper can be adequately described in terms of the reference point model. But how can we be sure that the referential choices in the examples, which are presumably the result of careful revision and editing, are not just a literary or stylistic artifact? The claim of the model presented here is much broader: that is, the reference point model represents a cognitive mechanism, which also underlies referential choices in spontaneous, non-literary narratives. As a further empirical validation of the reference point model, this section presents an analysis of discourse reference in an elicited *Dutch* narrative text.

Van Vliet (2008) reports the experimental elicitation and analysis of a corpus of Dutch discourse production data: participants were asked to produce a written narrative *online*, on the basis of a series of 25 pictures presented to them. The pictures, on which the narratives were based, portrayed a children's story featuring a single protagonist. The corpus

story presented below illustrates how consecutive references to the protagonist can be accounted for in terms of the factors described in this paper:¹¹

(32) Elicited Dutch Narrative:¹²

1. Enkele dagen voor de zomervakantie beslist **Maartje** een grote draak in elkaar te knutselen.
A few days before the summer holidays Maartje decides to build herself a large dragon.
2. Vol goede moed begint **ze** eraan en \emptyset wil door niemand gestoord worden.
In good spirits she starts on it and \emptyset does not want to be disturbed by anyone.
3. **Ze** haalt **haar** gereedschap boven en \emptyset knutselt ijverig voort.
She gets the tools from upstairs and \emptyset potters on diligently.
4. Na urenlang zwoegen, hoeft **ze** **haar** kunstwerk enkel nog te schilderen.
After toiling for hours, all she has left to do is paint her work of art.
5. Eindelijk is **haar** werk af.
Finally, her work is done.
6. Na enkele dagen is de vakantie begonnen en **Maartje** neemt haar draak overall mee.
After a few days the holidays begin and Maartje takes the dragon with her everywhere.
7. **Ze** heeft er niet beter op gevonden om de burens te laten schrikken.
She has not found anything better to do than scaring the neighbours.
8. sMorgens vroeg laat **ze** de draak verschijnen voor het raam van de buurman.
Early in the morning she makes the dragon appear in front of the neighbour's window.
9. Die lag nog lekker te soezen en \emptyset schrok zich natuurlijk te pletter.
He was still happily snoozing, so naturally \emptyset is scared out of his wits.
10. **Maartje** loopt razendsnel weg met de draak onder de arm.
Maartje walks away quickly with the dragon under one arm.
11. Enkele dagen later loopt **Maartje** alweer met de draak over straat.
A few days later Maartje is walking in the street with her dragon again.
12. **Ze** beslist nog maar eens een stunt uit te halen.
She decides to pull another prank.
13. Deze keer stopt **ze** de draak in de schoorsteen van een buur.
This time she shoves the dragon down a neighbour's chimney.
14. Voorbijgangers blijven staan om het kunstwerk te bewonderen.
Passers-by stop to admire the work of art.
15. Zodra de mensen uit het zicht verdwenen zijn, haalt **Maartje** de draak uit de schoorsteen en \emptyset holt snel weer weg!
As soon as they have disappeared from view, Maartje takes the dragon out of the chimney and \emptyset quickly runs off again!
16. De volgende dag gaat **Maartje** op zoek naar een nieuw avontuur.
The next day Maartje goes in search of a new adventure.
17. **Ze** trekt het bos in maar \emptyset loopt verloren.
She goes into the woods but \emptyset gets lost.

11. The quantitative analysis of the collected corpus of Dutch narratives, which focuses on consecutive references to the protagonist throughout the description of these 25 pictures, confirms a statistically significant effect of most of the factors described in this paper.

12. The numbers 1 through to 25 correspond to the numbers of the elicitation pictures.

18. Ze beslist een boom in te klimmen opdat iemand haar zou zien.
She decides to climb a tree so someone will be able to see her.
19. Vanop een grote afstand kan je de draak zien die boven de bomen uitsteekt.
From far away you can see the dragon sticking out above the treetops.
20. Maartje is zodanig moe dat ze boven in de boom in slaap valt.
Maartje is so tired that she falls asleep up in the tree.
21. De volgende ochtend wordt ze wakker en Ø merkt dat ze nog steeds in de boom zit.
The next morning she wakes up and Ø notices she is still in the tree.
22. Een eind verderop staan mensen de draak te bezichtigen.
Further on people stand looking at the dragon.
23. Ze besluiten om van dichterbij te gaan kijken en Ø vinden Maartje.
They decide to take a closer look and Ø find Maartje.
24. Maartje keert samen met haar bewonderaars terug naar huis.
Maartje returns home together with her admirers.
25. Ze wordt feestelijk onthaald en iedereen juicht Maartje en de draak toe!
She gets a warm welcome and everyone cheers on Maartje and the dragon!

References to the protagonist in this narrative display a number of factors discussed earlier in this paper. First, the episode boundaries in (6), (11) and (16) are all accompanied by repetition of the proper noun.¹³ The only episode transition in which pronominal reference to the protagonist is continued is sentence (21). This is possibly caused by the tendency to avoid the repetition of proper nouns in consecutive sentences. A further, tentative, explanation is that in the description of the protagonist *waking up*, the sentence is construed from the subjective POV of the protagonist, representing the protagonist as *conceptualizer*. (Note however that, in the description of *falling asleep*, in the previous sentence, the subjectivity factor does not override the factors distance and competition, cf. below.) Other instances in which the status of the protagonist as conceptualizer is accompanied by the use of pronouns are sentences (7), (12) and (18). Another relevant factor is the appearance of other, competing character references, and coincidentally, increased textual distance between references to the protagonist: the repeated proper nouns in sentences (10), (15) and (23) are triggered by referent competition as well as referential distance. A similar case is sentence (19), which is construed from a different point of view, and can also be said to involve implied or impersonal reference. Resumed reference to the protagonist in (20) occurs through a repeated proper noun. Lastly, the influence of salience *within* the clause can be observed in sentences (23) and (25), in which the protagonist functions as main clause object (rather than subject), and is referred to by a proper noun.¹⁴

13. The episode transitions were implemented as visual and textual cues in the picture stimuli for the narration task. The status of pictures 6, 11, 16 and 21 as new episodes was confirmed in a pretest.

14. Note that the use of null subjects is restricted to the second conjunct of coordinated main clauses. As in English, the use of null subjects is grammatically highly restricted in Dutch, and occurs almost exclusively in coordinated sentences.

7. Summary

This paper presents a reference point approach to *reference maintenance* in narratives. It is shown that Van Hoek's (1997) sentence level analysis, involving the interaction between referent accessibility and the strength of conceptual connections between coreferential nominals, also accounts for referential patterns at the level of (narrative) discourse. Reference point / dominion organization in narratives is influenced by (i) inherent topicality; (ii) competing referents; (iii) referential distance; (iv) conceptual connectivity; and (v) (embedded) subjectivity. A mental space approach to narrative representation – in terms of referential dominions – provides the partitioning and embedding of context that is needed to support an accessibility-based account of referential expressions in various usage situations. It is proposed that reference point organization forms part of attention framing, and as such contributes to the construal of referent salience and discourse connectivity.

References

- Ariel, M. 1988. Referring and accessibility. *Journal of Linguistics* 24: 65–87.
- Ariel, M. 1990. *Accessing Noun-Phrase Antecedents*. London: Routledge.
- Bolinger, D. 1979. Pronouns in discourse. In *Syntax and Semantics: Discourse and Syntax*, T. Givón (ed.), 289–309. New York: Academic Press.
- Chafe, W. L. 1976. Givenness, contrastiveness, definiteness, subjects, topics and point of view. In *Subject and Topic*, C. N. Li (ed.), 25–55. New York: Academic Press.
- Chafe, W. L. 1980. The deployment of consciousness in the production of a narrative. In *The Pear Stories: Cognitive, Cultural and Linguistic Aspects of Narrative Production*, W. L. Chafe (ed.). Norwood, N.J.: Ablex.
- Chafe, W. L. 1987. Cognitive constraints on information flow. In *Coherence and grounding in discourse*, R. S. Tomlin (ed.), 21–52. Amsterdam: Benjamins.
- Chafe, W. L. 1994. *Discourse, Consciousness, and Time. The Flow and Displacement of Conscious Experience in Speaking and Writing*. Chicago: Chicago University Press.
- Clancy, P. 1980. Referential choice in English and Japanese narrative discourse. In *The Pear Stories: Cognitive, Cultural and Linguistic Aspects of Narrative Production*, W. L. Chafe (ed.), 127–202. Norwood, NJ: Ablex.
- Deane, P. 1991. Limits to attention: a cognitive theory of island constraints. *Cognitive Linguistics* 2: 1–63.
- Dijk, T. A. v. 1982. Episodes as units of discourse analysis. In *Analyzing Discourse: Text and Talk*, D. Tannen (ed.), 177–195. Washington: Georgetown University Press.
- Emmott, C. 2003. Constructing social space: sociocognitive factors in the interpretation of character relations. In *Narrative Theory and the Cognitive Sciences*, D. Herman (ed.), 295–321. Stanford: CSLI Publications.
- Fauconnier, G. 1994. *Mental Spaces: Aspects of Meaning Construction in Natural Language*. Cambridge: Cambridge University Press.
- Fox, B. 1987. Anaphora in popular written English narratives. In *Coherence and Grounding in Discourse*, R. S. Tomlin (ed.), 57–74. Amsterdam: John Benjamins.
- Givón, T. 1983. Topic continuity in discourse: an introduction. In *Topic Continuity in Discourse: A Quantitative Language Study*, T. Givón (ed.), 1–41. Amsterdam: John Benjamins.

- Goldberg, A. 1995. *Constructions: A Construction Grammar Approach to Argument Structure*. Chicago: University of Chicago Press.
- Gordon, P. C., Grosz, B. J. & Gilliom, L. A. 1993. Pronouns, names, and centering of attention in discourse. *Cognitive Science* 17: 311–347.
- Grosz, B. J., Joshi, A. K. & Weinstein, S. 1995. Centering: a framework for modelling the local coherence of discourse. *Computational Linguistics* 21(2): 203–225.
- Karmiloff-Smith, A. 1981. The grammatical marking of thematic structure in the development of language production. In *The Child's Construction of Language*, W. Deutsch (ed.). London: Academic Press.
- Keenan, E. L. & Comrie, B. 1977. Noun phrase accessibility and Universal Grammar. *Linguistic Inquiry* 8: 63–99.
- Kuno, S. 1987. *Functional Syntax: Anaphora, Discourse and Empathy*. Chicago: Chicago University Press.
- Langacker, R. W. 1987. *Foundations of Cognitive Grammar* (Volume I). Stanford: Stanford University Press.
- Langacker, R. W. 1990. Subjectification. *Cognitive Linguistics* 1 (1): 5–38.
- Langacker, R. W. 1991. *Foundations of Cognitive Grammar* (Volume II). Stanford: Stanford University Press.
- Langacker, R. W. 2001. Discourse in Cognitive Grammar. *Cognitive Linguistics* 12 (2): 143–188.
- Lasnik, H. 1989. *Essays on Anaphora*. Dordrecht: Kluwer.
- Maes, A. 2001. A cognitive view of evolutive referents. In *Les référents évolutifs entre linguistique et philosophie (Recherches Linguistiques, no. 24)*, W. de Mulder & Co. Schneedecker (eds), 133–162. Paris: Klincksieck.
- Maes, A. & Noordman, L. G. M. 1995. Demonstrative nominal anaphors: a case of non-identificational markedness. *Linguistics* 33 (2): 255–282.
- Marslen-Wilson, W., Levy, E. & Komisarjevsky Tyler, L. 1982. Producing interpretable discourse: the establishment and maintenance of reference. In *Speech, Place, and Action*, R. J. Jarvella & W. Klein (eds), 339–379. Chichester: John Wiley and Sons.
- Sanders, J. 1994. *Perspective in Narrative Discourse*. PhD dissertation, Tilburg University.
- Sanders, J. & Redeker, G. 1996. Perspective and the representation of speech and thought in narrative discourse. In *Spaces, Worlds, and Grammar*, G. Fauconnier & E. Sweetser (eds), 290–317. Chicago: The University of Chicago Press.
- Schilperoord, J. 1996. *It's About Time. Temporal Aspects of Cognitive Processes in Text Production*. Amsterdam/Atlanta: Rodopi.
- Tomlin, R. S. 1987. Linguistic reflections of cognitive events. In *Coherence and Grounding in Discourse*, R. S. Tomlin (ed.), 455–480. Amsterdam: John Benjamins.
- Turner, M. 1996. *The Literary Mind. The Origins of Thought and Language*. Oxford: Oxford University Press.
- Van Hoek, K. 1995. Conceptual reference points: a cognitive grammar account of pronominal anaphora constraints. *Language* 71: 310–340.
- Van Hoek, K. 1997. *Anaphora and Conceptual Structure*. Chicago: The University of Chicago Press.
- Van Vliet, S. 2008. *Proper Nouns and Pronouns. The Production of Referring Expressions in Written Narrative Discourse*. PhD dissertation, Tilburg University.
- Vonk, W., Hustinx, L. & Simons, W. 1992. The use of referential expressions in structuring discourse. *Language and Cognitive Processes* 7 (3/4): 301–333.
- Werth, P. 1999. *Text Worlds: Representing Conceptual Space in Discourse*. London: Longman.
- Zwaan, R. A., Langston, M. C. & Graesser, A. C. 1995. The construction of situation models in narrative comprehension: An event-indexing model. *Psychological Science* 6: 292–297.
- Zwaan, R. A. & Radvansky, G. A. 1998. Situation models in language comprehension and memory. *Psychological Bulletin* 123 (2): 162–185.

Literary references

- Atwood, Margaret. 2003. *Oryx and Crake*. London: Virago (MA.OC).
- Cheever, John. The Country Husband, in: *The Stories of John Cheever*. 1980. New York: Ballantine (JC.CH).
- Dahl, Roald. 2001 [1961]. *James and the Giant Peach*. London: Penguin (RD.JGP).
- Dahl, Roald. 2001 [1988]. *Matilda*. London: Penguin (RD.M).
- Franzen, Jonathan. 2001. *The Corrections*. London: Fourth Estate (JF.TC).
- Grisham, John. 1993. *The Client*. London: Arrow (JG.TC).
- Hollinghurst, Alan. 2004. *The Line of Beauty*. London: Picador (AH.LB).
- McEwan, Ian. 2005. *Saturday*. London: Virago (IME.S).

The dream as blend in David Lynch's *Mulholland Drive*

Johanna Rubba

More than perhaps any other contemporary director, Lynch draws upon dream experience as a primal wellspring of his creative energy. Dreams and dreaming suffuse every moment of his approach to filmmaking. (Bulkeley 2003)

1. Introduction

[Dreams] are more likely to be the same types of figurative thinking that produce metaphor, metonymy, conceptual blending, and irony in waking life... (Domhoff 2001)

Dreams can be seen as self-contained mini-worlds of conceptual projections from our experience of waking life. Yet, while the stuff of dreams is taken from our conceptual structure, dreams have their own logic, revising, distorting, and defying reality. This makes them a compelling object for study within Fauconnier and Turner's conceptual blending (CB) framework (2002). Lakoff (1993, 1997) has made a case for interpreting dreams in terms of conceptual metaphors; CB provides tools for other correspondences between dream worlds and waking worlds. Its emphasis on emergent structure in the blended space is especially suited to understanding the peculiar logic of dreams.

This paper applies CB to David Lynch's film *Mulholland Drive*. My interpretation of the film's plot is that the first two-thirds of its running time is the protagonist's dream, an attempt to repair a life that has gone horrifically wrong. This section reviews folk and expert theories of dreaming; recent research indicates that mid-twentieth-century contestation of any psychological value of dreams was premature. I also describe how dreams are blends, and conclude with remarks on the centrality of dreams to the cinema of David Lynch.

1.1 Dream analysis

Humans have likely been interpreting dreams since they were able to interpret, that is, create meaning. How provocative dreams are – they deceive the dreamer by mimicking reality so convincingly; they can deliver to the dreamer extreme fear, grief, or joy equally powerfully, yet seemingly arbitrarily; and their origins are mysterious. Although they

appear in the dreamer's mind, the dreamer's lack of control over them invites speculation that they come from somewhere else – from the gods, from departed ancestors, from the devil, from the spirit world. Through the ages, they have been believed to give advice or warnings or predict the future, often through symbolism.

In the late-nineteenth and early-twentieth centuries, psychoanalytic theories introduced a new source of dreams, internal to the dreamer yet not less mysterious or inaccessible than the gods or the dead – the unconscious. Dreams were coded messages from the unconscious, requiring a new sort of shaman, the analyst, to help decode the message through free association.

A popularized, diluted version of psychoanalytic theories of dreaming is a common folk theory today, according to which dreams serve psychological purposes – conflict resolution, wish fulfillment, escape, or expressing repressed emotions.

Scientific study of the brain and sleep aided by body and brain-scanning technology has challenged psychological and spiritualistic theories of dreaming. First, dreaming was tied to REM (rapid eye movement) sleep, then to random firing in the pons, where no higher-order cognition occurs. The limbic system was also found to be active during dreaming, accounting for their emotional impact without appeal to psychological needs. More recent study, however, has revealed dreaming during non-REM sleep and activity in the rational pre-frontal cortex. “[T]here are probably multiple sources for the level of cortical activation that makes [dreaming] possible” (Vogel 1978, quoted in Domhoff 2001).

These developments have given rise to various new psychological theories of dreaming, e.g. the “reverse-learning” theory, which proposes that dreaming is a “brain dump”, unloading excess information at the end of the day (rather mundane, compared to the drama of raging, repressed urges). The 1970s activation-synthesis model proposed that the cortex attempts to build a coherent story out of the brain stem's random firings. Psychoanalytic theories are being revived, but other dream research contests that dreams cloak the unconscious. Domhoff (2001) cites numerous studies which show that “dreams are a reasonable simulation of the real world, and they usually concern everyday issues” and find “regularities relating to age, gender, and culture, and that a person's waking conceptions and concerns can be predicted [from] their dream reports”. Researchers at Harvard Medical School and the University of Maryland are exploring the value of dreams in problem-solving and mental-health therapy (Kantrowitz and Springen 2006).

Steen (1998) discusses yet another possible purpose dreaming serves. When the muscular paralysis accompanying dreaming is suppressed in lower mammals, they act out survival skills such as fighting and running, suggesting that dreaming is a kind of training. Fetal humans experience a great deal of REM, but it is doubtful that they need a *dream-work* to mask an unconscious Oedipus complex. Perhaps they are “practicing” instinctual behaviors they will need after birth. If dreaming trains physical behavior, perhaps it also trains the higher cognition peculiar to humans – mental and social survival skills. Might nightmares be practice for coping with extreme fear? Maybe dreams are a release-valve for socially regulated emotions such as anger and lust, whose expression would risk the social fabric if we acted them out. Research shows that waking mental “practice” – imagining the execution of a perfect golf swing, for example – improves physical performance. Perhaps dream survival training is similar.




Figure 1. Blend cartoon

1.2 Dreams as blends

This recent research indicates that we deploy our usual tools of cognition to make sensible stories out of neurological activity. Cortical processing of dreams may be just another example of the mind's drive to create narratives (Turner 1996). "Dreaming is our own storytelling time", says Rosalind Cartwright (quoted in Kantrowitz and Springen 2006).

The New Yorker cartoon (Figure 1) illustrates the folk theory of symbolic dreaming nicely, at the same time being an informal blending diagram: the thought balloon is the blended space, the dream; the "cast of characters" shows mappings between figures in the dream and persons in the dreamer's life. The nature of the characters (monster, stranger) leads us to infer that the dream expresses alienation from members of the dreamer's family.

Conceptual blending theory provides constructs for plotting the structure and dynamics of a dream: correspondences between dream elements and elements of conceptual structure; importation of schemas/frames into the dream space; and, most importantly, the emergent structure of the blend: the self-serving rules by which the dream unfolds.

Unlike in waking blends, "When you sleep, you don't control your dreams" (Lynch, in Rodley 2005: 15). Here lies the source of many theories of dreaming: what motivates the particular story that the cortex weaves from the neurological activity of the dream?

While we do not have conscious access to blending processes, they emerge from waking cognition. Their purposes are clear: efficient and effective reasoning. Less mundane blends, e.g. the Skiing Waiter (Fauconnier and Turner 2002: 21–22) and the Dinosaur (Fauconnier and Turner 2002: 93–96) show that blends are consciously constructed for didactic purposes. Dreams, however, are dictated to us by a freewheeling unconscious; only training allows us to control our dreams ("lucid dreaming"). What could the purposes of dream-blending be? Perhaps we do not have to abandon our folk theory – perhaps some dreams support psychological survival.

If dreams are blends, then interpreting a dream amounts to Unpacking¹ it: the meaning of the dream lies in the nature of the Vital Relations (VRs) within and between it and its inputs.

If one already has the entire network active, then running the blend gives inferences and consequences for the rest of the network. But if the entire network has not been built yet [or if parts of it are forgotten or not active], then the blend does good work in prompting for those activations. (Fauconnier and Turner 2002: 332)

Here, there is a big difference between waking and dreaming blends. Waking blends provide information from which inferences can be made. In the Skiing Waiter, for example, the SKIING frame is prompted by the context, a skiing lesson; the instructor uses the word *waiter*, which invokes RESTAURANT. In the common blend image of pre-human primates walking in line, turning into *Homo sapiens*, we are cued to compress the Time interval and Disanalogous identities by being informed of the communicative purpose of the image; for instance, there might be a date under each figure. We can also use background knowledge about evolution to infer what the image means.

Dreams often lack clear-cut prompts to activate the conceptual integration network's cross-space relations. In fact, the purpose of dream theorization is to find out what the rest of the network is. Are the inputs from the randomly-firing pons? The unconscious? The limbic system? Is the cortex striving to integrate signals from other brain areas into a comprehensible blend? What are the cross-space relations? A common dream is that one's teeth are falling out. We can map by identity from ourselves to ourselves, and we have a frame for teeth falling out, but what motivates the Integration of these two in the dream/blend? What inferences is the blend supposed to produce? Humans have puzzled over this since before Pharaoh.

Unpacking *Mulholland Drive* is challenging because it opens inside of a dream (see below). Vital relations between the dream and the diegesis aren't plainly cued (what does a young man in a diner have to do with a woman who has survived a car accident?) The inputs aren't made clear until the final third of the film, but even then Lynch deliberately obscures the VRs by having numerous Disanalogies and Changes: between actors and names in the movie's first vs. second parts, along with all of the quasi-duplicates: money, key, limo ride, etc.

Dreams are Decoupled, at least in terms of their effectuality in the real world (as long as the muscles remain paralyzed). This Decoupling is convenient if we have psychological purposes for dreaming: we can murder our boss in our dream, and thereby vent our anger safely. But the decoupling can be disappointing: Fido the lost dog may return home in your dream, but will be lost again when you wake up. Dreams are often counterfactual as well – for a dream to repair adverse circumstances requires this; the typical impossibilities of dreams (such as humans flying or animals talking) are also counterfactual and prompt efforts to make sense of the dream via symbolic or other interpretations.

1. Words capitalized in mid-sentence are to be understood as technical terms of CB theory.

1.3 Dreaming and David Lynch

As noted in the introduction to this paper, dreaming is central to Lynch's filmmaking. Bulkeley traces two facets of dreaming in many of his films: "dreams as revelations of truth, especially the truth of our deepest passions", as "giving voice to the primal wishes, fears, hopes, and aspirations of the human soul", and simultaneously their "sense of relentless, agonizing [sic] uncertainty about what is real and what is illusion" (2003:8). Dream plots suit several of Lynch's preoccupations: subjective identity; "the basic existential question of what is real" (Bulkeley 2003:4); and "how irrelevant it is to attempt to determine the precise level of consciousness at which we are existing" (Doniger 2001, cited in Bulkeley 2003:4). His reality-defying works are "enticing invitations to explore experiential realms beyond the boundaries of ordinary rational consciousness and personal identity" (Bulkeley 2003:5).

Lynch uses a kind of daydreaming in composing his films: "[w]aking dreams are the ones that are important, the ones that come when I'm quietly sitting in a chair, gently letting my mind wander ... I like to dive into a dream world that I've made or discovered, a world I choose" (in Rodley 2005: 15). He recounts how he developed *Mulholland Drive* from the name "Mulholland Drive": "[w]hen you say some words, pictures form ... [t]his makes me dream, and these images are like magnets and they pull other ideas to them" (in Rodley 2005:270).

Bulkeley (2003) lays out three major ways Lynch uses dreaming: (1) as a narrative structuring device; (2) characters' dreams; and (3) discussion or mention of dreams in a film. All three are found early in *Mulholland Drive*. The first two-thirds of the film's narrative is the protagonist's dream; there is a dream within this dream, along with mentions of dreams by characters.

The clue to the protagonist's dream occurs about two minutes after the film begins. The camera zooms steadily into a crimson-cased bed pillow. We hear labored breathing; then the image fades to black. The dream sequence begins at this moment; the next thing we see is the "Mulholland Drive" street sign. We do not see the dreamer's face or body; the camera is seeing through her eyes in a subjective treatment of the dreamer as focaliser. In this way, the dreamer's perceptions are the spectator's perceptions. Lynch deliberately takes the viewer into the dreamer's mind so that s/he can experience the raw emotion as directly as the character: "right there is the power of cinema ... the dreamer has bought [the fiction of the dream] 1,000 per cent ... It's so unique and powerful to that person. But with sounds and situations and time you could get much closer to putting that together for somebody else with a film" (Rodley 2005: 15).

This statement is the clue to understanding why Lynch filmed *Mulholland Drive* as he did. Through violation of numerous conventional film techniques (to be discussed in detail below), Lynch reproduces in the spectator the profound disorientation, drama, and emotion of this protagonist's dreams; in the last third of the film, the uncued dive into the protagonist's flashbacks maintains the disorientation. This protagonist is a woman suffering extreme despair, grief, and remorse, still mixed with heartbreak and anger over how others have treated her. She wakes from a powerful dream to face (via the signal that the hired murder of her lover has been carried out) the reality that was undone in her dream.

The texture and color of the film, Angelo Badalamenti's score – suggestive of darkness even in the opening sequence of dancers jitterbugging to swing music – and the melodrama of the dream's "script" convey the protagonist's state of mind. The spectator feels dread, without really knowing why. As the addled psychic character Louise declares, "Someone is in trouble! Something bad is happening!" The spectator is left to infer the rest.

Lynch maintains the disorientation through to the roll of the final credits: the protagonist dies, but we still are left to ourselves to figure out what was real and what imagined in her life (indeed, some viewers suggest that *none* of the film is diegesis – that every scene is in some sort of mental space of a character, even the final suicide and subsequent shots).

One must be perspicacious enough to notice that the bed in which the dreamer awakes is clothed in the same crimson sheets as in the opening sequence. This bracketing is a major clue to the fact that the events in between are a dream. More evidence for the dream is presented in Section 2 below.

Lynch's other typical uses of dream motifs continue. During the first quarter hour, a beautiful woman survives a murder attempt by a man driving her in a limousine on Mulholland Drive; she stumbles down a canyon to Sunset Boulevard and eventually sneaks into an apartment on Havenhurst Drive. (All of this is geographically accurate except for the Havenhurst address.) There, she falls asleep under the table in a breakfast nook. The scene then shifts in a shock cut to a young man in a diner called Winkie's who is recounting a *nightmare* he has had about an evil man behind the diner. Then the scene cuts equally abruptly back to the sleeping woman; the implication is that the diner scene is her *dream*. Soon a young woman named Betty Elms arrives at the apartment, which is her aunt's. She has come to Hollywood to try her luck in the movies; in telling her story to the trespassing woman, she says "I mean I just came here from Deep River, Ontario – now I'm in this – *dream place*" (emphasis added).

Various metaphors built on dreaming suffuse the film. Several critics see it as a parable about a familiar schema of American popular culture: the *dream* of success in Hollywood's *dream factory*. The protagonist of this schema is a young woman, an ideal *dream-girl* sought after by movie-makers. The melodrama of the *Mulholland Drive* dream scenes and dialogue simulate the script of classic Hollywood movies telling such a story. Watts' scripted, slight overacting of the Betty role, with the perky, innocent enthusiasm of the dream-girl-to-be, evokes these sometimes-corny films. Her unexpectedly successful audition is lifted right out of such plots. Perhaps not incidentally, Lynch at one time was to direct a film of the life of Marilyn Monroe, the paragon of the dream-girl (and a dream gone wrong) if ever there was one (Rodley 2005: 268).

The dreams in *Mulholland Drive* are of course not real dreams; this is a fictional story. They are, however, *realistic*. Dreams usually feature people important in the dreamer's life, but may also include acquaintances or briefly-encountered strangers who take significant roles in the dream. Events are often absurd or illogical. Changes in real-life events and characters fit the purposes of the dream: a dead loved one is alive again; a lover who has dumped you returns, repentant. The *Mulholland Drive* dream possesses these traits.

Most interestingly with regard to *Mulholland Drive*, there are hints that the true facts are quite different from those of the dream and sometimes irrupt into our dreams, dragging us back to reality. The resurrected loved one may dig a hole and lie down in it; the

repentant lover may suddenly have to leave for a distant place. Several such hints occur in the dream of Diane Selwyn, *Mulholland Drive*'s protagonist: the evil bum behind Winkie's, in which location, in the diegetic story, the protagonist paid to have her unfaithful lover killed. And there is the visit to Betty by the psychic Louise, who declares that "Betty" is not Betty's name (Betty is Diane's re-creation of herself in the dream).

Bulkeley sees Lynchian films as examples of "the power of dreaming to relieve people's suffering by imagining different and better lives for themselves" (2003). Lynch's own comments about some of his films indicate that he believes altered states of consciousness can satisfy psychological needs: they are "the mind tricking itself in order to save itself from having to deal with the un-dealable" (Lynch, in Rodley 2005: 289). This, he says, was the core idea of his film *Lost Highway* (which also features a protagonist who turns into a different person). The core of *Mulholland Drive*, he says, is "a love story" (quoted in Rodley 2005: 289), with the mind-trick coming in to serve that theme.

In the same passage, he notes the cost of repression: "The mind is such a friend to us when it shuts off certain things. But there's a price to pay for shutting it off. It can fester ... It's a beautiful place, but it can also be pitch dark" (Rodley 2005: 289). Hence the nagging intrusion of Diane's reality into her dream.

The remainder of this paper, after a summary of *Mulholland Drive*, uses CB to derive how the dream is an attempt to "deal with the un-dealable". Understanding the film as a Conceptual Integration Network enables us to Unpack the motivation for the dream.

2. *Mulholland Drive*²

Mulholland Drive, like so many of Lynch's films, is multi-layered, extraordinarily subtle and complex, subject to analysis at many levels. The story line is embedded in a rich texture of allusions to the movie industry itself and features some of Lynch's favorite themes and images: Los Angeles in the Golden Age of Hollywood; red curtains; a Mystery Man; doppelgängers; love, sex, and violence; psychotic breaks; a plot that unwinds more like a Möbius strip than a spool of film.

At the first viewing, *Mulholland Drive* appears incoherent. Perhaps because the film was released widely as a feature and shown not just in art houses, its incoherence was roundly criticized. Ironically, the plot of the film actually unfolds in normal chronological order: in narratology terms, the discourse, or presentation of the content, does not differ from the story, which is the content being narrated. Fabula and sjužhet are one. The narrative viewpoint from which the story is filmed obscures this.

Two main characteristics of the film cause much of the confusion: shock cuts between seemingly unrelated plot lines, and the sudden change in the protagonist's identity from

2. Release date: October 2001; Running time: 2 hours 27 minutes; Director: David Lynch; Screenplay: Joyce Eliason and David Lynch; Producers: Pierre Edelman, Mary Sweeney, Alain Sarde, Neal Edelstein, Michael Polaire, Tony Krantz; Music: Angelo Badalamenti; Cinematography: Peter Deming; Major characters: Betty Elms/ Diane Selwyn: Naomi Watts; Camilla Rhodes/Rita: Laura Elena Harring; Adam Kesher: Justin Theroux.

Betty Elms, promising cinema ingénue, to Diane Selwyn, an unkempt and haggard depressive. This accompanies a major plot change from a promising love story to a bitter tale of betrayal. Abundant absurdities compound viewers' puzzlement: an apartment tenant with a fighting kangaroo; hyperbolically angry and demanding Mafiosi who glower, shout, and spit espresso; a miniature, ghostly elderly couple squeezing under the front door of an apartment, growing to human size, and making comical pretend-scary gestures which nonetheless literally frighten a young woman to death.

But there are dozens of clues to putting the puzzle back together; Lynch himself notes that "[a]ll the threads in *Mulholland Drive* are tied up" (in Rodley 2005: 287). These clues take the form of people and things that appear twice, once in the first two-thirds and again, usually in a different guise, in the last third. Among these are a black limousine driving along Mulholland Drive; Diane's telephone ringing; Diane's grandparents; a chenille bathrobe; a thirties-era tiled kitchen; a blue key; a wad of cash; Winkie's diner; actresses' head shots; even lines of dialogue: "we don't stop here"; "this is the girl". The problem with these duplications is that they are not presented in traditional match cuts, in which the props remain, but their context, and sometimes features, change in an immediate transition (think for example, of a typical shot showing, say, a telephone ringing within a character's dream; the shot then transforms to the character's actual phone ringing in its "real" location in the dreamer's dwelling). Lynch decompresses his match cuts. For example, about seventeen minutes into the film, we see a series of telephone calls transmitting the message "the girl is still missing". The last shot in this sequence is of a particular telephone on a small table in a dark room, lit by a small table lamp. Over an hour and a half later, the exact shot is repeated; this time, we are able to identify the phone as that of the protagonist, Diane Selwyn.

Nearly all the characters seen in the first third also show up in the last third, again in decompressed match cuts all the more confusing because their identities change: Coco, the apartment manager in the dream, becomes the mother of Adam, the director; Rita, the beautiful amnesiac, becomes Camilla Rhodes, a less-than-savory love object; sinister figures from part one, such as the Cowboy and a Mafioso, are briefly-glimpsed guests at a party. The function of these connections as clues is disrupted because of the distance in time between the shots and the stark differences in the circumstances of their second appearance. What the viewer doesn't realize is that the clues are appearing in reverse order – we don't understand the relations between the doubles until the last third of the film. This film is a case in which the writer/director's cinematic choices are "so unconventional that the result is a puzzlingly opaque account of what turns out to be a fairly ordinary and familiar event" (Semino and Swindlehurst 1996: 1), in this case, a woman remaking her failed life in a dream.

Rodley writes:

Conventional film narrative, with its demand for logic and legibility, is ... of little interest to Lynch ... Insecurity, estrangement, and lack of orientation and balance are sometimes so acute in Lynchland that the question becomes one of whether it is ever possible to feel 'at home'.
(Rodley 2005: x–xi)

Despite constant requests, Lynch resists explaining or interpreting his films. He desires that people experience his films at the same intuitive level at which he creates them and that they sort out what they mean to *them*.

I think people know what *Mulholland Drive* is to them but they don't trust it ... I love people analyzing it but they don't need me to help them out. That's the beautiful thing, to figure things out as a detective. Telling them robs them of the joy of thinking it through and *feeling* it through and coming to a conclusion.

(Lynch, quoted in Rodley 2005: 289; emphasis in original)

My conclusion is that *Mulholland Drive* is the “evil twin” of the small-town-girl-makes-it-big-in-Hollywood story: Diane Selwyn, naïve native of “Deep River”, Ontario, inspired to seek movie stardom after winning a jitterbug contest, uses an inheritance to travel to Hollywood. But her movie career is a non-starter, and she is betrayed and humiliated by her beautiful yet vapid and scheming lover, Camilla Rhodes. Rage drives her to hire out the murder of Camilla. She receives the signal that the hit has been carried out (a blue key). Overwhelmed by guilt, grief, and fear (detectives are after her), she goes to bed and has a dream, which occupies the film's first two hours. She wakes from the dream; unable to overcome her misery, she suffers delusions and kills herself.

As noted above, the narrative viewpoint of the first two-thirds of the film obscures the fact that it depicts a dream. There is no clue in the film's opening as to whether we are in the diegesis or viewing events subjectively through the mind's eye of a focaliser. Lynch is especially clever in how he sets this up: we first see a colorful collage of couples in period costume jitterbugging to 1950s rock'n'roll; shining heads of a young woman and an elderly couple (Diane smiling into applause at the jitterbug contest; Diane and her grandparents) appear, jitter, fade, and reappear; there is a partial fade to black, then the camera shifts from a blurry view of unidentifiable objects – the view of the floor as one would see it lying face-down on the edge of the bed. The focus then clears to the view of the bed made up in crimson sheets. The breathing we hear is Diane's; the next blackness comes with the closing of her eyes. Apparently, the dancing couples are in a dream; she awakens to blurry vision which clears as she lies down and falls back to sleep. From the blackness, the “Mulholland Drive” street sign appears; then the opening credits appear over the nighttime view of the limousine on Mulholland Drive. The viewer is misled into believing that those first images are intended to frame the “real” story.

From the first moment, the camera stands in for the protagonist in a visual stream of consciousness, a one-hundred-percent subjective treatment. We are in Diane's mind, seeing through her eyes. We do not move to an objective treatment until two hours have passed.

2.1 Order of events

The order of events of the film's diegesis is as follows:

1. Diane's arrival and failure in Hollywood and her affair with Camilla.
2. Camilla breaks off her relationship with Diane.

3. Camilla's engagement party; Diane's humiliation.
4. Diane hires the hit man.
5. The hit is carried out.
6. Diane goes to bed.
7. She has her dream, which picks up between steps 4 and 5, just in time to save Camilla from being killed.
8. She wakes from the dream and comes back to reality.
9. She sees the key.
10. She fantasizes herself experiencing the return of Camilla and flashes back to 1 through 4.
11. A knock at the door rouses her from her flashbacks; she hallucinates the elderly couple (who may in reality be the two detectives her neighbor warned her about), and shoots herself.

These events occur in the film in the order 6, 7, 8, 9, 10, 11. 1 through to 5 become known via 10.

Diane's dream goes back in time to a moment just before the murder, but from then on, even the chronology of the dream-story is normal. The seeming discontinuities arise from normal filmic *decompression* of time: several of the dream's subplots (Adam's misfortunes; the hit man) run parallel to the main plot (the Betty-Rita story), with the movie scenes switching from one to another. This is a convention of cinema, as the only alternative is split-screen presentation. Unlike conventional films, however, *Mulholland Drive* is poor in obvious transitional cues (e.g. bridging shots, dissolves, sequential match cuts, subtitles – *meanwhile, back at the ranch* – or voice-over narration). What cues Lynch does offer usually serve to obfuscate further rather than clarify. For example, the transition from the dream to Diane's waking up is achieved through cuts between fades to black – about as clear a bridge as Lynch delivers in this film – between the hallways of the dream-apartment and the hallways of Diane's diegetic apartment; then the camera moves into her bedroom. To keep things mysterious, the Cowboy appears and tells Diane to wake up, but the figure we see on the bed is a corpse that featured in the dream. With more cuts separated by total blackness, the corpse becomes the living, sleeping Diane, on the same bed, and in the same pose, but dressed differently. To complicate things even further, when Diane shoots herself at the very end of the film, she winds up in the same pose, on the same bed – hence the speculation that we are not yet out of the subjective viewpoint even at this juncture.

Other typical movie devices are used in deliberately confusing ways – for instance, after Diane wakes and sees the blue key, we see her standing at the kitchen sink in what looks like an objective shot. She looks to her left and her expression changes to joyful incredulity; shoulders heaving, body trembling, she utters with extreme emotion, "Camilla! You've come back!" What she sees is revealed in a shot/reverse shot; we switch to an eyeline shot from the sink, and Camilla, alive and stunningly beautiful as ever, is standing there; a small smile and a blink of her eyes suggest contrition. This suggests that the view of Diane at the sink was Camilla's subjective point of view. The eyeline shot reverts to Diane, and, in a masterful stroke of acting and makeup, the lovely Diane becomes a washed-out, sniveling supplicant. We see her dawning realization of the horror of her act, as her expression

morphs into abject terror, going to the edge of monstrosity. When the shot returns to the POV from the sink to where Camilla was standing, however, we see a calmer, disgusted-looking Diane looking back at where she seemed to have been standing a moment before. Without further clues, the viewer must puzzle out how Diane could look at a supposedly dead Camilla, who then seems to vanish, to be replaced by Diane looking at herself.

The implication is that, in this scene, Diane is imagining herself witnessing Camilla's return to their love affair, an event that would have prevented the ensuing murder. She is disgusted by how she would have reacted in that situation – groveling, self-abasing gratitude. The scene shows Diane's realization of how much of her dignity she had sacrificed in falling prey to Camilla's casual exploitation of her affections. If not for the grimness of her crime, this would be a glimmer of hope: gaining an objective view of herself as a first step towards reclaiming her self-worth. In this one shot, Watts conveys a woman who has moved a long way towards maturity, but, alas, she is fated not to recover. As she moves towards the couch with her cup of coffee, her flashbacks to the scenes of her degradation begin; then the knock at the door penetrates her semi-conscious state, with the subsequent hallucination and suicide.

The remainder of this paper presents a CB analysis of Diane's dream. Because of the complexity of *Mulholland Drive's* blends, I will devote little discussion to generic spaces, except for the initial layout of the blend. There is not space here to discuss the main conceptual integration network completely, or to bring in the numerous smaller correspondences, such as the money, the blue key, the Club Silencio show, and key aspects such as the bum at Winkie's and the Cowboy.

3. Why does Diane dream?

As noted above, I am proposing that Diane's dream is an attempt to revise reality. As such, we expect cross-space VRs between Inputs and the Blend of Change and Disanalogy that combine with Cause–Effect, Role, Property; interesting angles on Identity and Intentionality; Compression and Decompression could also play important roles. An optimal dream would Integrate these connections tightly and would reflect and perhaps adjust scales in domains such as emotion, success/failure, and power.

If Diane dreams to remake reality, there should be no misfortune: her dream counterpart should be, and is, successful in both career and love; Camilla should not be, and is not, jaded and patronizing. However, Adam should not only reverse his decision to choose Camilla over Diane for the lead role in his movie; he should remain a starmaker. There should be no murder for hire, no Mafiosi, no wad of cash in Rita's purse. Yet evil carries over to the dream. This will be explicated below.

Diane's goals in dreaming are as follows:

1. To reconstruct her acting career as successful;
2. To reverse Camilla's rejection of her love;
3. To erase grief over Camilla's death;
4. To change Camilla into a good person;

5. To reverse the power Camilla and Adam wielded over her;
6. To release her from guilt over Camilla's murder; and
7. To exact revenge on those who abused and humiliated her (including Camilla).

Diane's dream is a tightly Integrated, double-scope, counterfactual conceptual integration network comprising four spaces: Input 1 is the film's reality, Diane Selwyn's misfortunes; Input 2 is Diane's wish space – how she wishes things had happened (inferred from the film and general knowledge); the Blend (the dream) and the Generic space. There are two primary organizing frames: SMALL-TOWN GIRL MAKES IT BIG IN HOLLYWOOD and LOVE RELATIONSHIPS.³ Within these, particular subschemas are recruited to optimize the dream's effectiveness: SMALL-TOWN GIRL subsumes MOVIE PRODUCTION; MOVIE PRODUCTION, in turn, hosts MAFIA. Mafia-related movie corruption was part of the Hollywood Golden Age Lynch alludes to, and movies featuring Mafia goings-on were also a popular genre in that era. MAFIA then enables EXTORTION and MURDER FOR HIRE, which in turn enables GIRL DETECTIVE. LOVE RELATIONSHIPS includes BETRAYAL as a possible, albeit undesirable, outcome of falling in love. How the frames are distributed across the spaces is shown in Table 1.

The table shows that some frames appear only in the dream. Diane brings them in in order to facilitate the resolution of her various conflicts, as will be elaborated below. Table 2 shows which changes between reality and the dream serve particular goals.

Table 3 shows the main VRs and structural differences that turn 1a, 2a, etc. into 1b, 2b, etc.

The remainder of this section will consider viewpoint and Vital Relations between the characters. Section 4 will lay out how mappings across frames and completion of the Blend satisfy all of Diane's goals.

Table 1. Distribution of frames across the *Mulholland Drive* conceptual integration network

Frame	Input 1: Reality	Input 2: Wish	Blend: Dream	Generic space:
LOVE RELATIONSHIPS	x	x	x	x
BETRAYAL	x	–	x	–
SMALL-TOWN GIRL MAKES IT BIG IN HOLLYWOOD	x	x	x	x
MOVIE PRODUCTION	x	x	x	x
MAFIA	–	–	x	–
MURDER FOR HIRE	x	–	x	–
GIRL DETECTIVE	–	–	x	–

3. I use SMALL CAPS for the name of frames/schemas. Names of component spaces and Vital Relations are capitalized: Blend; Disanalogy.

Table 2. Reality/dream differences and goals they serve

a. Reality b. Dream	Goal
1a. Diane is talentless, mousy, vengeful 1b. <i>Betty is talented, adventurous, smart, confident, happy</i>	1, 2
2a. Diane's Hollywood career fails 2b. <i>Betty's Hollywood career is promising</i>	1
3a. Camilla uses, betrays, humiliates Diane 3b. <i>Amnesiac Camilla ("Rita") falls in love with Betty</i>	2, 4
4a. Adam chooses Camilla over Diane for lead role 4b. <i>Female lead forced on Adam by Mafiosi</i> 4b. <i>Adam is mesmerized by Betty</i>	1, 7 5, 7
5a. Adam steals Camilla's love from Diane 5b. <i>Adam's wife betrays him with the pool man</i>	7
6a. Unscrupulous Camilla has the upper hand 6b. <i>Camilla/"Rita" is helpless, childlike, fearful</i> 6b. <i>Betty takes control of solving "Rita"'s mystery</i>	2, 4, 5, 7 5
7a. Diane hires a hit man to murder Camilla 7b. <i>Someone else has hired out the hit</i>	6
8a. Camilla is killed 8b. <i>Camilla survives thanks to a car accident</i>	3, 6

Table 3. Vital Relations differentiating dream from reality

1a-b: <i>Disanalogous Identity</i>
2a-b: <i>Disanalogous Cause/Effect</i>
3a-b: <i>Disanalogous Properties > Disanalogous Cause/Effect</i>
4a-b: <i>Disanalogous Role/Value, inverse power (topology)</i>
5a-b: <i>Analogous Roles and emotional force dynamics; Disanalogous Role/Value</i>
6a-b: <i>Change in Properties > inverse power (force); inverse emotional force</i>
7a-b: <i>Disanalogous Role/Value</i>
8a-b: <i>Unique Cause-Effect: car accident causes survival; inverse Effect</i>

3.1 Viewpoint: Stream-of-consciousness vs. first-person/third-person compression

3.1.1 *Mental spaces, space-builders, and connectors in Mulholland Drive*

We interpret communication by keeping track of what mental space we are in at any given moment. Part of that tracking is understanding whose viewpoint the scene is being viewed from. Point of view has been a rich subject of investigation in cognitive linguistics as well as in film and literature studies, particularly narratology. Understanding whose point of view is framing each scene is crucial to unlocking the mysteries of *Mulholland Drive*.

The business of making meaning out of fiction is a business of building mental spaces cued by the work (therefore by its creator), and deploying mappings over those spaces.

A fictional film-story is a mental space within which the filmmaker builds a world and its unfolding events; the audience plays along with idea that the “suspension of disbelief that we all perform before entering into a fictional world entails an acceptance of a story’s diegesis” (Felluga 2003). Thus the framing discourse space of fiction is the creator of the work addressing the work to readers/viewers, and all are complicit in understanding the story as not true. (In documentary filmmaking, on the other hand, the world built is not fictional, but is understood metonymically as standing in for the real events, allowing the audience to respond somewhat as they would if they had witnessed them. The audience members are genuine – albeit second-hand – observers, most often with a narrator explicitly addressing them).

Our amazing ability to interpret works of fiction rests ultimately in our innate ability to *pretend*, something young children can do long before they can pronounce the word. Within the discourse space of a filmmaker presenting to an audience, cultural conventions allow interpretation of the scenario as communication (the filmmaker’s message) or as “sheer entertainment”. We make complex mappings from actors to characters; we build mental spaces as dictated by the setting and time period portrayed in the film; we build a vast number of mental spaces as we follow the conversations of the characters while projecting their emotions and thoughts as cued by their acting. Conventions such as flashbacks and flashforwards, moves from subjective to objective POVs, identity/change mappings from, e.g., characters as children (often played by different actors), as youths, and transformed into old age by cutting-edge makeup techniques, are all blending operations that allow us to suspend the rules of real time, real people, and real places. Fantasy worlds such as those of *Planet of the Apes*, *Dune*, and the distant future in *Star Trek* pile on more blends.

Thus, acting is space-building: thanks to our ability to construct complex blends and distinguish fiction from reality, and to distinguish watching a film from actually experiencing the portrayed events,⁴ we conform to the convention of using actors’ behavior and lines to build mental spaces in which we imagine the characters’ responses, thought processes, emotions, and so forth. An example from *Mulholland Drive* is the extremely short scene which I propose as being straight film narration, featuring the eyeline shot scene at the kitchen sink. When the camera reverses to view the space where Camilla stood a moment before, we see Diane, clear-eyed and restored to her usual attractiveness, looking disgustedly at the place where her horrific imagined self stood. We take her expression as a cue to build a space in which we project onto her the sorts of inner reactions that we associate, in real life, with such facial expressions: my interpretation above that this moment signals her realization of how she sacrificed her dignity to Camilla’s manipulations, and her disgust at how deeply she had debased herself.

It is impossible in the space here to give anything near a full exposition of the mental spaces constructed and cued in *Mulholland Drive*, but an examination of the primary mental spaces and how they are cued or not cued reveals Lynch’s techniques of portraying life realistically and confusingly at the same time. Depriving viewers of conventional

4. Recall that, in the early days of cinema, audiences rushed from their seats to escape things like “oncoming trains” onscreen.

film space-builders and connectors that allow them to infer the correct point of view is a major part of his technique. As noted above, the film's discourse and story follow normal chronological order: the framing space of the film is Diane's story, and the time line begins with her dozing and dreaming, moves through her waking and flashbacks, and ends with her hallucination and suicide. The dream also tells its story in normal chronological sequence, from Camilla surviving the murder attempt to Betty and Rita's return to the apartment from *Silencio*. The dream alternates conventionally between the Kesher/Mafia plot and the Betty/Rita plot.

The sources of *Mulholland Drive* viewers' confusion are the schema violations of the dream's events coupled with Lynch's deliberate refusal to use explicit filmic space-builders and cross-space connectors most of the time. Schema violations are quite ordinary in dreams, and it is quite conventional in film for the story to alternate between plots that we are to understand as simultaneously unfolding, and that are viewed from multiple viewpoints. Lynch's violation of film conventions is how he achieves the "[i]nsecurity, estrangement, and lack of orientation and balance" Rodley speaks of (2005: x-xi).

There are few primary mental spaces in *Mulholland Drive* (see Figure 2). Most of the film unfolds in Diane's mind: first, the dream; after she awakes, the brief incident in which she imagines herself reacting to Camilla's return (see 2.1 above). Then come the flashbacks in which she recalls the crisis points of her breakup with Camilla and her final hallucination of the "menacing" old couple. We spend but a few minutes in the conventional discourse space of film: a story narrated to the audience by the filmmaker through the medium of cinema. (1) describes the main mental spaces of the film.

(1)

- *Diane's mind*: the dream (nearly 2 hours)
- *Diane's mind*: Rita's dreams within Diane's dream
- *Film's Reality space*: Diane awakes; neighbor scene (about 3 minutes)
- *Diane's mind*: Diane imagines herself reacting to Camilla's return
- *Film's Reality space*: Diane makes coffee and carries the mug of coffee to the couch (about 1.5 minutes)
- *Diane's mind*: flashbacks - Camilla ends the affair; Diane's humiliation on the movie set; angry scene at Diane's door; masturbation scene; Diane goes to the engagement party; Diane paying the hit man; tiny elderly couple come out of paper bag at bum's camp (about 18 minutes)
- *Film's Reality space*: Diane on the couch in the evening, staring at the blue key, growing terrified; seeming to lose consciousness (less than a minute)
- *Diane's mind*: her final hallucination
- *Film's Reality space?* Diane's suicide; her dead body on her bed; smoke; bum's face; city at night; shining images of Betty and Rita dressed as they were at *Silencio*; fade to *Silencio's* red curtain and the Blue-haired Lady on the balcony, who utters, "Silencio". Credits roll. (about 1.5 minutes)

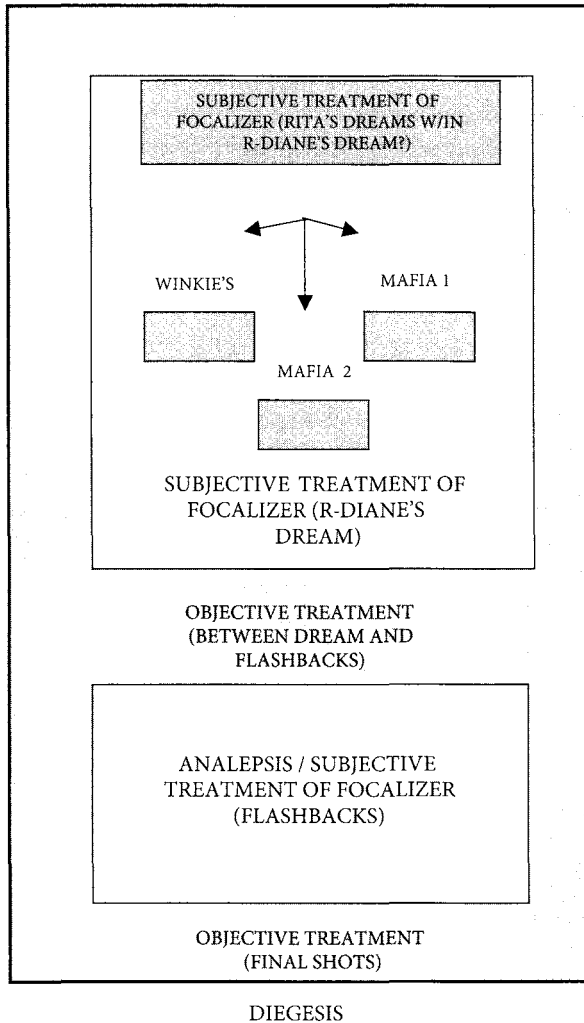


Figure 2. Mental Spaces in *Mulholland Drive*

Lynch gives us only these space-builders:

1. The camera zoom into the pillow, and the fade to black, at the beginning of the film. These are, however, seen subjectively, from within Diane's mind.
2. There are a few space-builders within the dream, which will be discussed below.
3. The transformation from the corpse on the bed to the sleeping Diane.

The remaining space-builders have to be inferred by the viewer, who, unfortunately, begins the film "lost in space" and must wait a very long time for space-building cues, which then prove not to be particularly helpful.

Consider a few alternatives: Lynch could have begun the film with the zoom shot into the pillow, the breathing and the fade to black, instead of the montage of jitterbug-

gers against a flat, brightly colored background. This would have been a clearer hint of the dream to come. To be even more conventional, he could have opened with a shot of Diane sleeping on her bed, then zoomed in to a close-up of her face, then faded to black. Instead, he begins the film inside of Diane's doze/dream and uses a blurred shot to move from this dream to her partial waking, return to sleep and the longer dream. The blur, and the mismatch between qualities of the jitterbuggers, the floor/bed, and the black night on Mulholland Drive are confusing without an instruction to the viewer to build the space of Diane's mind. With such a cue, a viewer would quite easily accept the strangeness of the two opening shots (dancers and blurred floor). On top of this, the verisimilitude of the Mulholland Drive limousine ride, attempted murder, and car accident intensify the confusion, for it seems as if we have moved from something unreal to something real, when, in fact, we are still inside of Diane's mind.

These scenes are, however, quite realistic: often, when dozing, we think or see things that don't make sense. Often, our vision is blurred when we awake from the doze, and clears as we move about. When we close our eyes to sleep, everything goes black. We are not aware of the falseness of a dream; as Lynch says, we buy it "one thousand per cent". Our mind does not record the time between our falling asleep and beginning to dream, and it does not make a distinction between reality and the dream.

As noted above, the transition which takes us out of the dream to Diane's waking up and answering the door, instead of being a clear transition, is merely another mystery (the dark hallways; the cuts of the corpse morphing into a living Diane). The physical-resemblance Identity connector between Betty and this Diane does not make sense to the viewer. Again, a more conventional device would have been to replicate the shot of a sleeping Diane from the start of the film, thus bracketing the dream unambiguously, and to have the loud knocking awake her. But, as with other transitions, it is, while artful, not unrealistic. At that point in the dream, Betty and Rita are about to solve the mystery of the blue key. The key's meaning is what Diane cannot deal with: she has murdered the woman she loves. Instead of allowing this into the dream, she has Betty disappear, then Rita, then the blue box; that "it was all a dream" is, within the dream itself, keyed by Aunt Ruth's checking the bedroom (because she heard a noise?) but seeing nothing on the floor where the box fell. This, like the limousine shot early in the film, once again leads the viewer to believe that the world of Aunt Ruth is real. However, when Aunt Ruth exits, her dark hallway begins to transform into r-Diane's, and the latter's dream connection with the corpse on the bed (who carries the name *Diane Selwyn* in the dream) provides the final return to reality.

The end of the dream may be prompted by the knock on the door – we often take a while to rouse from sleep when hearing a loud noise, and our knowledge of our real setting often overlaps with the final scenes of a dream: Diane hears the knocking, and her mind begins the transition to consciousness by transforming the hallway of Aunt Ruth's apartment to her own (with a clever visual pun,⁵ a trope Lynch seems to like: the knocking is loud enough "to wake the dead"). Such mixes of dream and reality are common. For

5. Another such pun in the film is the blue key, appearing as a futuristic sculpture in the dream, and as a mundane household key in the reality space: it is the "key" to the awful truth of Rita's identity and Diane's crime.

instance, dreamers often incorporate a real sound that is bringing them to consciousness into their dreams: a real ringing telephone becomes a dream classroom bell or fire alarm; a barking dog suddenly appears as a character in the dream.

As noted above, most of the transitions between mental spaces in the dream and between the dream and the film's reality space are uncued. They are quite dramatically uncued: these changes of scene are accomplished via shock cuts so abrupt and extreme that they disorient a viewer's actual perception. Often, the time of day, mood, and color schemes contrast profoundly. Yet a clue is sometimes there.

The first of these disorienting shocks occurs early in the film: the sequence from the limousine ride to Camilla/Rita's falling asleep in Aunt Ruth's apartment is very conventionally shot, realistic, and follows a logical order. As soon as she rests her head on her purse, however, there is a split-second switch to a screen-filling close-up of the Winkie's diner sign in bright daylight. Though this is an establishing shot, the abrupt change of lighting and location weaken it. There is no fade to black, no zoom into a close-up of Camilla's face, or of her eyes slowly closing. We are then treated to the bizarre story of Dan and the bum "who's doing it". After Dan collapses, there is an abrupt switch back to Camilla sleeping. This shot is brief, and then there is another split-second switch to another radically different scene: a telephone exchange between Mr. Roque and a large man sitting in an ornate room announcing that "the girl is still missing". There is no apparent connection between the Dan story and the Roque story. (2) shows the order of this sequence.

(2)

1. Camilla/Rita falls asleep under Aunt Ruth's breakfast-nook table after surviving the attempted murder and car accident.
2. Winkie's Diner, where Dan tells his story and is shocked/killed by the sight of the evil bum.
3. Camilla/Rita sleeping under the table.
4. Roque and unknown man telephoning.
5. Betty arrives at Los Angeles International Airport.

The shots of the sleeping Camilla bracket the Winkie's story, suggesting, as does the crimson bed linen early and late in the film, that it is a dream, but the viewer's lack of background knowledge about Camilla, and the total unrelatedness of the Winkie's story to the attempted murder, lay a heavy interpretive burden on the viewer. To compound the problem, there is no close-bracket shot of Camilla sleeping after the telephone scene. Instead, the story makes another abrupt switch to Betty's arrival at the Los Angeles airport.

A similar sequence occurs after Betty arrives at the apartment and discovers "Rita". Camilla/Rita falls asleep, then there are shock cuts to two more Mafia sequences; again, the first Mafia scene is bracketed by shots of her sleeping, but the second is not. Instead we get a conventional film space-builder, a fade to black. But instead of returning to the sleeping Camilla/Rita, we are transported to the shabby neighborhood inhabited by the hit men. There are no more scenes of Camilla/Rita sleeping until after the love scene, about an hour later in the film.

The hit-men scene is cued by an establishing shot zooming into the exterior of the building in which one hit man's office is located, but, again, the shot is more disorienting than helpful.

As if this weren't enough, the Mafia/Kesher story and hit men plot continue after Rita is awake and well, and the Mafia/Kesher plot merges with the Betty/Rita story. And when we discover the significance of Winkie's and Dan (Diane paid the hit man there before the murder, and glimpsed Dan at the cash register), there is no logic in Camilla's dreaming about an event she couldn't have known about. Our strategy of using the bracketed shots of Camilla sleeping to conclude that the Dan, Mafia and hit men plots are her dream becomes untenable – how can the dream continue when she is awake?

Another very disorienting transition occurs after Diane has awakened and made herself a mug of coffee. We are in third-person viewpoint mode; we can interpret this scene as standard film narration. Diane is walking up to her couch from behind it; the camera tracks her from close behind, with a view only of her midsection. When she gets close enough to look over the back of the couch, we are greeted by the sight of r-Camilla, reclining topless on the couch. The camera then cuts to the opposite angle, and we see, not only Camilla reclining, but Diane climbing over the back of the couch, now dressed in jeans and also topless. Instead of a cup of coffee, she is carrying a drink, which she sets down on the coffee table.

Seductive music is playing; the mood is that of a languorous love scene, contrasting starkly with the disgust and depression of the preceding scenes. But the scene devolves into angry conflict: Camilla announces the end of their affair. There is no cue to tell us we are now in Diane's mind, flashing back to the turning points of her story. Conventional film techniques such as cutting from a tracking shot to a subjective POV, then to a third-person view become disorienting because of the abrupt changes in the features of the people and the mood.

All of these uncued mental spaces are realistic. If the Winkie's and Mafia scenes are Camilla/Rita's dreams, their abrupt appearance mimics our lack of awareness that we have fallen asleep and are dreaming; the bizarre story of the bum is realistic when understood as a typical dream schema violation. As to the couch scene, it is quite plausible that someone who has just committed a heinous act out of jealousy and rage would recall vividly the events that led up to the act. Tricking us yet again, Lynch actually provides these scenes to explain the correspondences between the dream and the diegesis. It takes a good deal of ex-post-facto comparison and inference to solve the puzzle, for the connectors between the dream and Diane's reality are also confusing (see Sections 3.2 and following on *Vital Relations*).

Lynch's violation of conventions for signaling point of view thus leaves viewers confused and insecure about their ability to distinguish what is "real" (in the film story) from what is not, which is precisely what he wants.

Such techniques are not unknown in written fiction, of course. Alain Robbe-Grillet's ironically-named "objectivist" style uses exactly the same trick as does Lynch: abrupt, unexplained scene changes. In Robbe-Grillet's case, the cues are usually paragraph breaks, but not much more. In the opening paragraphs of his novel *Jealousy*, for example, he moves abruptly from a detailed description of the sunlight/shadow play on the veranda of a house to a description of a woman entering a room with a window that gives onto that veranda. Her movements about the room are described, including the fact that she does *not* look towards this window. At the end of the paragraph, the woman turns her head,

and a new paragraph opens which describes in detail the balustrade of the veranda and the gardens beyond. The narrating viewpoint is, in fact, completely subjective: a person on the veranda is assessing what can be seen from inside the room, and watches the woman as long as she is not looking towards the window; the moment she turns such that she can see through the window, the watcher turns away. Thus, if the woman sees this person, she sees someone innocently observing the garden, not someone watching her.

Certain scenes in the story occur again and again, with slight changes; there is much speculation on the part of the narrator as to what exactly happened and what various characters meant by what they said. The reader has to infer that we are hearing the voice of an obsessive character, going over and over certain events in memory, trying to gauge their significance with respect to the relationships among the various characters. We are hearing the narrator's fears, suspicions, and obsessive ruminations, but the lack of overt space-builders leaves this discovery up to us.

3.2.1 *Viewpoint compression*

Lynch also manipulates the role of the camera in establishing viewpoint. Three conventional viewpoint techniques for film are important in analyzing *Mulholland Drive*. The first is the third-person observer viewpoint: the camera is a "bystander", seeing things as a person watching, but not participating, would. This is especially apparent when the camera "sees" events that one or more characters couldn't be aware of. In Alfred Hitchcock's *Psycho*, for example, the audience sees Norman Bates peeping at Marion Crane through a hole in the wall between the motel office and her room. The audience's fear and suspicion of the Bates character grow, while Crane remains ignorant of and vulnerable to the danger she is in. This is a standard maneuver of horror films and murder mysteries.

Then there are different sorts of third-person/first-person compression: The camera mediates between the audience and the internal experience of the characters; the two viewpoints are compressed. This is found in thousands of shots in film and television in which the camera flips back and forth between a third-person view of a character and that character's subjective view (and this view is where the compression takes place). In police and detective stories, for instance, we are given third-person shots of detectives sitting in a car or van surveilling a suspect, alternating with shots of the setting and people they are watching, from their angle of view.

Another type is the eyeline shot, in which the camera alternates between the subjective POVs of two conversing characters. In each shot, we apprehend two viewpoints: that of the addressee watching the speaker, for whose eyes the camera is standing in. We also interpret the inner workings of the speaker's mind via the actor's lines and acting cues such as a sarcastic smile or eyes widening in fear. We surmise both characters' thought processes based on our understanding of the speaker's lines and facial expressions and our knowledge of the story so far. This is a dual compression of our viewpoint with that of each of the characters, respectively: we simultaneously take in what the unseen observer is seeing (and thinking) and project the thoughts and feelings of the speaker.

The third viewpoint technique is the totally subjective stream-of-consciousness viewpoint, in which the camera *is* the eye/mind of a character: we experience the perception of the character unmediated (suspending, in accordance with the filmmaker-viewer con-

tract, our knowledge of the “fiction” discourse context, of course). Returning to *Psycho*, the camera moves from an external view of Bates peeping at his view, to a subjective POV, framed by the peephole. We see Marion Crane as she moves about her motel room, oblivious of being watched. As we watch the film, we experience simultaneously the (normal for the viewer) discomfort of watching someone who thinks she is alone, and (by projection into Bates's mind) the repulsive fascination of the peeper.

These compressions are devices so conventional that we hardly notice them: in written third-person omniscient narration, the narrator's outside view is collapsed with the characters' interior view, resulting in lines such as *She thought she would never be happy again* or *She hoped he would come back*. We are simultaneously inside and outside of the character's mind; but, usually, space-builders such as *she thought* and cross-space connectors such as *she* cue us to the compression. As noted above, novelists like Robbe-Grillet violate the convention of providing such space-builders and connectors.

As we have seen, there is a mix of these techniques throughout *Mulholland Drive*. There are many third-person shots of both Betty and r-Diane, mixed in with subjective treatments, and Lynch uses film conventions to obscure rather than clarify which is which. In dreams and remembering, we sometimes see the dream events completely subjectively, as we see the world in waking life. Sometimes, however, we see ourselves as a character, sometimes looking like ourselves, sometimes as someone/something else (like the “little boy” in the *New Yorker* cartoon). Lynch deliberately keeps us guessing, playing viewpoint games such as that of the sink scene.

3.2 Who's who: Vital Relations among the characters

[B]lending is a powerful and supple instrument for creating and disintegrating identity.
(Fauconnier and Turner 2002:95)

“We'll pretend to be someone else.” (Naomi Watts as Betty, *Mulholland Drive*)

In folk theorizing about dreams, much symbolic value is derived from the nature of Vital Relations between the input spaces and the blend. The *New Yorker* cartoon in Figure 1 links role values from the dreamer's waking life (father, mother, brother) to Disanalogous roles/values found Uniquely in the blend: monster, stranger. These mappings supposedly reveal the dreamer's true feelings about the person.

One of the most fundamental human blending operations is the recognition of the identity of people and objects across time, space, and situation – we recognize our siblings or children and consider them to be the same people day after day. We recognize our car in the parking lot at the end of every work day. Our schematic knowledge can so strongly influence our perception of places and things that we might not see things that are there, as happened to an insurance claimant who stated, “Arriving home, I pulled into the wrong driveway and hit a tree I don't have” (Cross n.d.).

Lynch plays with this basic operation to dramatic effect: in the last third of the film, after Diane awakes from her dream, there are many connectors, both visual (physical resemblances; places and objects) and linguistic (names), between this and the first two thirds of the film. These should provide continuity, but go awry. The former Rita is Camilla

Table 4. Actress/role mappings

Actress	Diegesis	Dream
Naomi Watts	Diane Selwyn (r-Diane)	Betty Elms
Lyssie Powell	–	d-Diane
Laura Elena Harring	Camilla Rhodes (r-Camilla)	Rita
Melissa George	Woman who kisses r-Camilla at party	d-Camilla

Table 5. Vital Relations between characters across spaces

Input 1	1 > 2	1 > B	2 > B
Diane Selwyn	Id/Ch*: Diane Selwyn	Ds: Betty Elms	An: (someone like) Betty Elms
Camilla Rhodes	Id/Ch: Camilla Rhodes	Id/Ch: Rita	Id/Ch: Rita
Adam Kesher	Id/Ch: Adam Kesher	Id/Ch: Adam Kesher	Id/Ch: Adam Kesher
Woman who kisses Camilla at party	–	Camilla Rhodes	–

* Id = Identity; Ch = Change; Ds = Disanalogy

Rhodes, but does not at all resemble the Camilla Rhodes of the dream; the former Betty Elms is now Diane Selwyn, who was dead in the dream. Yet Adam Kesher remains Adam Kesher. Lynch's mixing of these usually reliable cues once again disorients the viewer.

There are Identity relations between a few important characters in Diane's waking life and the dream. Adam Kesher remains himself, as does Joe the hit man. Both Change to some degree – Adam in his Roles and fortunes, and Joe similarly. Dan at Winkie's is the same person in the dream and reality, but his story of the nightmare is Unique. The VRs between persons in the inputs and the Blend are shown in Tables 4–5.

In sum, Naomi Watts plays two roles in the film: the real Diane Selwyn, and Betty Elms, the “dream-girl”. Laura Elena Harring plays two roles as well: in the film's reality, she is Camilla Rhodes, the dark-eyed, dark-haired beauty; in the dream, she is Rita. Camilla and Rita are, technically, the same person in the world of the dream – Rita is Camilla with severe amnesia. Things become confusing when the two names, *Diane Selwyn* and *Camilla Rhodes*, turn up in the dream, but are played by different actresses: Lyssie Powell and Melissa George, respectively. It is not too difficult to manage Diane's transfer of her identity to Betty Elms, her ideal for herself, and her displacement of her real name, the label of failure, onto a completely different woman. More confusing is the seeming presence of two women named Camilla Rhodes in the dream, especially considering d-Camilla's appearance: she is a blue-eyed blonde. Moreover, at the engagement part in the last third of the film, d-Camilla kisses r-Camilla. This is a trick similar to those of M. C. Escher's famous drawings, e.g. the trompe l'oeil placement of pillars in *Belvedere*. The ladder's foot is inside the building, yet its top is outside. In the dream, Rita is Camilla, and yet Camilla is not Rita. The benefit to Diane of this conundrum will be explicated presently.

Under *Identity*, the essence of the person remains unchanged; the person remains the same person. In *Change*, the essence of the person carries over, but there may be changes in her/his state, roles, age, appearance, or even personality (Fauconnier and Turner

2002:93). *Disanalogy* entails a difference in essence: disanalogous entities across spaces are not the same thing/person (Fauconnier and Turner 2002:99).

Mulholland Drive's mappings between people are a prime cause of the confusion, and therefore complaints, of viewers and critics. Lynch is well-known for such mappings: in *Lost Highway*, "a complex cross-weave of parallel worlds and identities" (Rodley 2005:215), the protagonist changes into a completely different person, then, near the end of the film, changes back into himself. He even hears a message for himself from himself over a front-door intercom. Lynch's television series *Twin Peaks* features an actress who plays both the protagonist and her cousin. *Mulholland Drive* is another of Lynch's explorations of the subjectivity of identity; understanding how *Mulholland Drive*'s changes and disanalogies satisfy Diane's psychological needs is key to untangling the film.

The sections below will explain how the Vital Relations given in Table 5 serve the purposes of Diane's dream.

3.2.1 *Diane/Betty*

r-Diane seems to have two counterparts in the blend: Betty Elms and the dead woman that Betty and Rita find at Diane Selwyn's address. This is reflected in the table by the Disanalogy relation between r-Diane in Input 1 and both of the corresponding characters in the Blend. These are the most important Vital Relations of the film: r-Diane Selwyn delinks herself completely from her actual identity, transferring it wholesale to Betty Elms. It is not hard to imagine Diane saying to herself, "I wish I were someone else." Betty is that someone else: a new person, Unique to the Blend, in the role of SMALL-TOWN GIRL. She is unlike Diane in all but physique (both are played by Naomi Watts), a vital Disanalogy in Properties between Input 1 and the Blend. It is crucial to understand that this is not Change. Diane leaves her essence as a failure and murderer behind; Betty has traits that will bring her success, admiration, and love.

Two pieces of "physical evidence" support this Disanalogy. One is that the corpse which Betty and Rita find on the bed when they visit d-Diane's apartment is played by a different actress, Lyssie Powell.⁶ And, most crucially, Betty arrives in Los Angeles *after* the attempted hit and *after* d-Diane's death. This removes her from the murder plot, serving Goal 6: relieving r-Diane of her guilt. How could Betty, arriving after the attempted murder, possibly be involved? Why would she want to kill someone whom she hasn't even met yet? This is about as effective an alibi as one could wish for.

This Disanalogy is similar to Fred Madison's change to Pete Dayton in *Lost Highway*. Both employ the same strategy: "Fred Madison and Diane Selwyn are forced to adopt extreme measures to achieve the illusion of stability and happiness, creating more innocent parallel identities and worlds for themselves – dream scenarios in which events struggle to overcome the reality of mental collapse" (Rodley 2005:xi). The innocence is achieved by

6. The abbreviational prefixes *d-* and *r-* are used henceforward to distinguish between different characters with the same name in the dream vs. reality: *d-Diane* refers to the dead woman on the bed in the dream (played by Lyssie Powell). The waking Diane is *r-Diane*. *d-Camilla* is the woman forced on Adam by the Mafia to play the lead role in his new film; she is played by Melissa George. *r-Camilla* refers to the real Camilla, whom the real Diane has murdered. She is played by Laura Elena Harring.

Disanalogous Intentionality assignments: r-Diane's intention to kill r-Camilla is projected onto someone else (Mafiosi, as we shall see below). Betty can have pure Intentions to help Rita; r-Diane's simultaneous love and hatred of r-Camilla are decompressed: Betty has only the love, and the Mafiosi only the evil intent.

Furthermore, Betty has Properties that invert some of Input 1's topology of personal power and emotional force dynamics, as well as Cause and Effect: now a talented actress, she impresses a major casting agent with her first audition. Her charisma mesmerizes Adam as he is pretending to audition actresses for his movie. These are reversals of Cause-Effect VRs in Input 1. She also has intelligence, verve and daring, equipping her to attack Rita's mystery as the GIRL DETECTIVE. Bringing this frame into the Blend alludes to yet another American theme of yesteryear, as manifested in series like the *Nancy Drew* mysteries, made into movies during the 1930s. The girl detective is a hero – her mystery-solving serves the public good and helps individual people (here, Rita). AS GIRL DETECTIVE, Betty takes the clever precaution of using a public phone instead of Aunt Ruth's to call the police to inquire about the accident; she daringly climbs into d-Diane's apartment through an unlocked window; in the following scene at Aunt Ruth's, when Rita wants to disguise herself, Betty takes charge: "I know what you have to do ... let me do it".

Betty is trusting and altruistic as well, accepting Rita, a total stranger and trespasser dressed to the nines, with no ID and tens of thousands of dollars in her purse. She helps her regardless of any risk to herself. She even sacrifices a possible breakout role by leaving Adam's audition to help her.

These properties enable Betty to exert emotional force over Rita/Camilla: Rita, in her helpless state, allows Betty to make major decisions such as entering the apartment (illegally); she depends on her. Betty's altruism, talent, and down home congeniality abet Rita's attraction to her.

The name change from Diane to Betty has several meanings. It is symbolic of the wholesale change of Diane's personality and fate in the dream. She "receives" the name at a climactic moment in Input 1: Betty is a waitress who serves coffee to r-Diane when she is setting up r-Camilla's murder at Winkie's with Joe the hitman – the same scene in which she glimpses Dan at the diner's counter. Perhaps r-Diane is recalling that this is the moment that sealed her fate, a moment when she wishes she were someone else, somewhere else. Subconsciously taking the name *Betty* provides the release she desired at that moment, or later as her fear and remorse grew. *Betty* is also a name more characteristic of Hollywood's golden years and alludes to great movie stars of the time, such as Bette Davis.

It is safe for r-Diane to keep the name in the Blend, as long as it is attached to someone different onto whom r-Diane can project her failures and shortcomings. She also needs a Diane Selwyn in order to preserve the link between r-Camilla's real-life infractions and Rita, thereby achieving one part of goal 7: REVENGE against r-Camilla. r-Diane wants payback: r-Camilla's complicity in her humiliation and betrayal cannot go unpunished. So r-Diane keeps her own drive to self-immolation in the dream, so Rita can intuit the magnitude of hurt she (as r-Camilla) has done to r-Diane. In her amnesiac, regressive state, Rita is horrified by the dead woman, where in real life she might have been disgusted by the corpse, but perhaps not moved much beyond pity by the death of the insignificant Diane Selwyn.

The implied connection between Rita and d-Diane's death also contributes to Rita's fear, bolstered by two suited men apparently looking for her. Rita becomes so frightened that she feels she must disguise herself after discovering the body.

Scalar adjustments occur in the Blend as well. I remarked previously on the melodrama of the dream's events and dialogue; similarly, Betty's traits are a trifle idealized – how many would-be stars would be altruistic enough to skip out on the chance to audition for a major role, having made an instant, deep connection with the director? Wouldn't someone new to a big city hustle an injured trespasser to a hospital or to police? Betty's hypergoodness is the stuff of the wish world.

Rita's extreme reaction upon seeing the corpse is Analogous in nature and scale to r-Diane's horror at her murder of r-Camilla, but it doesn't stop there. Why does Rita feel *guilt* when she sees the dead woman? Why does she assume she is complicit in the death? The woman could have killed herself, or been killed, for any number of reasons. Once again, r-Diane is projecting – Rita's fear is similar in scale to her own fear of being punished, as is Rita's guilt. r-Diane turns the tables in an ultimate irony: she projects her own guilt over killing r-Camilla onto Rita, simultaneously salving her conscience and punishing r-Camilla's dream counterpart.

3.2.2 *Camilla/Rita*

The Vital Relations involving Camilla Rhodes may be even more challenging than the Diane/Betty mapping: we are to understand that Rita is indeed a Changed r-Camilla; yet, another woman named Camilla Rhodes is in the dream. Once again, r-Diane needs an illogical arrangement to serve her psychological goals: 2, requited love between her and r-Camilla; 4, a chance for r-Camilla to reform; 5, correction of the power asymmetry between her and r-Camilla, and 7, revenge against Adam and r-Camilla for the heartbreak and humiliation they dealt her.

The VRs between r-Camilla and Rita are Identity and Change. Her amnesia allows her to be Camilla and not-Camilla at the same time; she is the same person, but with changes in her outlook and reactions to events, and with a chance to start anew. r-Camilla and d-Camilla (actress Melissa George) are, however, linked by Disanalogy. In Diane's waking life, Melissa George plays a woman who kisses r-Camilla at the engagement party. As with the name *Diane Selwyn*, the dreaming Diane preserves the name *Camilla Rhodes* for purposes of revenge, this time against Adam (see below). Rita's amnesia conveniently makes the name available.

r-Diane despises Adam for two reasons: he chose r-Camilla over r-Diane for the lead role in *The Sylvia North Story*, and he wooed r-Camilla away from r-Diane. These two events destroyed r-Diane's hopes. In the dream, Adam is in for a dramatic comeuppance. The punishment for the first offense comes from the Blend's Unique Mr. Roque and the Mafiosi (perhaps under contract to Roque), who dictate Adam's choice of leading lady. This is a serious blow for an egotistical director. He gets a Camilla Rhodes, but not the one he wanted – in fact, he gets a woman who was messing around with his fiancée, probably behind his back. This is delivered by compression of r-Camilla's and d-Camilla's identities via the connector inherent in the name. At the same time, this is

a decompression, for r-Camilla and her name have become detached in the dream. The handling of *Diane Selwyn* is a similar decompression.

Rita's amnesia is essential to goals 2, 4, and 5. r-Camilla is jaded and opportunistic; through Identity accompanied by Change, Rita is fearful and helpless. The amnesia is provided by a Cause–Effect chain Unique to the dream: the car accident, which simultaneously saves Rita from death and erases her memory. She gets a chance to remake herself; losing her memory, she also loses her corrupt self (she cannot remember her *name*, a symbol as well as a fundamental aspect of her identity). Gone are her memory of a relationship with Adam, of flings with other women; of her affair with r-Diane; of the failed actress's story. Her timid, sweet new persona (Change in Properties, an Effect of the amnesia) turns her into a woman one can both love and control – and maybe forgive. This permits importation of an Analogous *REQUIRED LOVE* from Input 2 to the Blend to trump the *UNREQUIRED LOVE* of Input 1.

Rita's amnesia also helps satisfy needs 5 (power) and 7 (revenge). The power (interpersonal force dynamics) polarity is also reversed: Betty, r-Diane's Blend counterpart, has the upper hand. Helpless Rita is dependent on her: Betty shelters her, cares for her, and takes the lead in discovering her circumstances. She protects her from the meddlings of Aunt Ruth, Louise, and the apartment manager. Betty is a benevolent power, however, making her superior to those who used their power over her malevolently.

Goal 7 requires an element of serious danger for Camilla. As noted above, r-Diane both loves and hates Camilla. The dream enables her to both punish and redeem her: Camilla's slate is not swept completely clean. She has amnesia rather than simply being a different kind of person. The injury, the loss of memory, and the implications that she is guilty of some transgression and is being hunted cause her suffering. r-Diane eats her cake and has it: r-Camilla, as Rita, suffers for her offenses, while r-Diane, as Betty, is her hero.

As with the Diane–Betty name change, the Camilla–Rita name change is symbolic. Rita is a “cleansed” Camilla; perhaps the loss of her memory has restored some pure core of goodness in her – she is not only kind to Betty, but can fully measure the horror of her misdeeds implied by the dead woman, the cash, and the men looking for her.

These two name changes are significant in yet another way: they introduce a good–bad polarity. The real-life names are attached to undesirable women in the dream. d-Diane, the counterpart of r-Diane's detested self, is dead by violence; d-Camilla is not only the counterpart of a woman who kisses r-Camilla at the engagement party, but also, in the dream, the woman forced on Adam as his leading lady. Rita's recall of the name *Diane Selwyn* links her to the corpse and “transmits” whatever evil the corpse implies to her.

Similarly, in *Lost Highway*, protagonist Fred Madison (Bill Pullman), morphs in prison into a different man with a different name: Pete Dayton (Balthazar Getty). The astonishing change causes Dayton to be released from prison. As Betty frees r-Diane from responsibility for r-Camilla's death, Pete Dayton relieves Fred Madison of guilt for a heinous crime: murdering and dismembering his wife. In a feat of Freudian projection, each transfers guilt to a different person. This leaves their alter-egos clean candidates for goodness (seen in *Mulholland Drive's* Betty; Fred is not seen while Pete is in *Lost Highway*).

4. Mappings between roles and relations

The value of the Blend in *Mulholland Drive* is its Unique, emergent interactions and events, which repair r-Diane's misery and avenge her maltreatment. In this section, these will be presented by showing correspondences within and across structural frames given above in Table 1: SMALL-TOWN GIRL MAKES IT BIG IN HOLLYWOOD; MOVIE PRODUCTION; LOVE RELATIONSHIPS; MAFIA; and MURDER FOR HIRE.

4.1 Love relationships

Other things being equal, set up the blend and the inputs so that useful topology in the inputs and their outer-space relations is reflected by inner-space relations in the blend.

(Fauconnier and Turner 2002: 327)

The *Mulholland Drive* dream is designed to give tit-for-tat: those who abused r-Diane should experience powerlessness, humiliation, guilt, and betrayal. Therefore it is imperative that relations between roles⁷ in Input 1 map into the Blend, but it is also crucial to bring in those from Input 2. We therefore find a lot of relations in the Blend that are Identical or Analogous to those in the Inputs, and the tit-for-tat is delivered by changing Role/Value relations. Especially important is the topology of metaphorical force dynamics – interpersonal power. Table 6 shows the relevant Roles and their relations in the LOVE frame; Table 7 shows the cross-space Vital Relations among these in Diane's dream/blend.

As noted in Section 3, the dream/blend reassigns identities or imposes changes on the people from Input 1. This set of relations is one part of the dream which preserves the undesirable offenses and evils from reality that would be absent in the wish world (Input 2), but are needed to accomplish Diane's psychological goals. The Role/Value bindings and disanalogous connectors in the network place the "good guys" (r-Diane, reformed r-Camilla in the person of Rita) and the "bad guys" (Adam, bad r-Camilla) in their "correct" positions to change unrequited love to requited love (goal 2), and to give Adam a taste of his own medicine (goal 7) by finding out how it feels to be betrayed and cuckolded. Rita, as the "missing girl" replaced by d-Camilla, doesn't get to see Adam, forestalling any potential trigger of Rita's memory.

Table 6. Roles/relations in LOVE RELATIONSHIPS

LOVE RELATIONSHIPS		spaces: 1	2	B
a. REQUITED	<i>X loves Y and Y loves X</i>	–	x	x
b. UNREQUITED	<i>X loves Y and Y ~ love X</i>	x	–	–
c. BETRAYAL	<i>X betrays Y</i>	x	–	x
d. CUCKOLDING	<i>X cuckolds Y</i>	x	–	x

7. By "relations" here I mean not Vital Relations across spaces, but properties of elements within a frame, for instance the relations: unequal power and the extortionist's desire for money or whatever action is desired from the victim; the victim's fear of the extortionist, etc.

Table 7. Vital Relations of Roles/relations between spaces in LOVE RELATIONSHIPS

Relation	Pattern	I	I > 2	I > B	2 > B
REQUIRED	X loves Y, Y loves X	–	Ds: D loves C, C loves D	Ds: B loves R, R loves B	An: B loves R, R loves B
UNREQUIRED	X loves Y, Y ~ love X	D loves C, C ~ love D	–	Ds: R loves B Ds: A loves B?	Ds: A loves B?
BETRAYAL	X betrays Y	C betrays D	–	Ds: L betrays A	Ds: L betrays A
CUCKOLDING	X cuckolds Y	A cuckolds D	–	Ds: G cuckolds A	Ds: G cuckolds A

* Ds = Disanalogy; An = Analogy; – = no counterpart; C = r-Camilla; D = r-Diane; R = Rita; B = Betty; A = Adam; L = Lorraine (A's wife); G = Gene (pool man, co-respondent of L's betrayal)

Table 8. Mappings in the SMALL-TOWN GIRL MAKES IT BIG IN HOLLYWOOD frame

Roles/ Outcomes	Input 1 Reality	Input 2 Wish	Blend: Dream	VRs 1 > B	VRs 2 > B	VRs 1 > 2
<i>Small town Girl</i>	<i>Deep River, Ont.</i>	<i>Deep River, Ont.</i>	<i>Deep River, Ont.</i>	<i>Id</i>	<i>Id</i>	<i>Id</i>
<i>Outcome</i>	<i>Failure</i>	<i>Success</i>	<i>Success</i>	<i>Ds</i>	<i>Id</i>	<i>Ds</i>

Table 9. Mappings in the MOVIE PRODUCTION frame

Roles/ Outcomes	Input 1 Reality	Input 2 Wish	Blend: Dream	VRs 1 > B	VRs 2 > B	VRs 1 > 2
<i>Movie</i>	"Sylvia North Story"	"Sylvia North Story"	A's next hit	<i>Ds</i>	<i>Ds</i>	<i>Id</i>
<i>In control</i>	Adam	Adam	Roque/Mafia	<i>Ds</i>	<i>Ds</i>	<i>Id</i>
<i>Director</i>	Adam	Adam	Adam	<i>Id</i>	<i>Id</i>	<i>Id</i>
<i>Auditioners</i>	D, C	D, C	B, d-C, others	<i>Ds</i>	<i>Ds</i>	<i>Id</i>
<i>Lead actress</i>	C	D	d-C	<i>Ds</i>	<i>Ds</i>	<i>Ds</i>

4.2 Small-town girl makes it big in hollywood

Tables 8 and 9 show how a careful distribution of Identity and Disanalogy VRs satisfies r-Diane's goals of success and revenge.

r-Diane's new incarnation, Betty Elms, immediately attracts attention with her steamy audition, fulfilling r-Diane's wish for a chance at stardom. Her appearance on Adam's set provides a few more twists of the knife of his misfortune: his frustration at the forced choice of leading lady is compounded by the intriguing Betty's arrival – and prompt disappearance. This scene also delivers a dollop of satisfaction of goal 5 – r-Diane's wish to turn the tables on those who have power over her. Betty's charisma captivates the egotist who humiliated r-Diane.

4.3 Mafia

[To satisfy important blending principles] [i]n a double-scope network ... It is necessary to use a frame that has been developed specifically for the blend and that has central emergent structure. (Fauconnier and Turner 2002: 340)

We have seen that *SMALL-TOWN GIRL MAKES IT BIG IN HOLLYWOOD*, *LOVE RELATIONSHIPS*, and *MURDER FOR HIRE* are imported into the Blend from Input 1, with VRs set up so as to repair r-Diane's misfortunes. *MAFIA*, however, is Unique. This elaborate plotline seems excessive, if its only function is to Change the Value of the murder culprit. As it turns out, *MAFIA* tightly integrates Analogous topology and relationships from the inputs with the differences in scalarity and VRs we have seen above. Pattern completion of these elements in the major organizing frames of the Blend maximizes its effectiveness in relieving r-Diane's pain as she elaborates the Blend in her dream.

Only two global frames are needed: *SMALL-TOWN GIRL MAKES IT BIG IN HOLLYWOOD* and *LOVE RELATIONSHIPS*. *SMALL-TOWN GIRL* encompasses *MOVIE PRODUCTION*; within the latter, a *CORRUPTION* frame is activated in order to provide a plausible scenario for both Camilla's murder and Adam's misfortunes. *CORRUPTION* is specified as *MAFIA EXTORTION* (another part of Lynch's depiction of the seamy side of the industry, and another allusion to common Golden Age movie themes). Other elements of *MAFIA CORRUPTION* are payoffs and witness targeting, hence *MURDER FOR HIRE*. These together produce the mystery of Rita, whose encounter with Betty, in turn, provides Betty with the opportunity to be *GIRL DETECTIVE*. *GIRL DETECTIVE* and *SMALL-TOWN GIRL* can then be tied into *LOVE RELATIONSHIPS*. *LOVE RELATIONSHIPS*, in turn, hosts *BETRAYAL*. And we're off! These relationships are shown in Figure 3.

Figure 3 shows the hierarchical structure of the various frames that structure the Blend. Each arrow points to a frame or frame element embedded in it. The arrows originate in a larger frame which contains the frames or elements at which they terminate. The various arrow paths trace the subplots of the dream and show how each subplot hooks into the others.

We may use formulae to trace each path. In (3) to (8) below, each frame or frame element is shown in small caps. The '>' symbol indicates which frame/element belongs to the one that precedes it, reading from left to right. A slash (/) indicates a Role within the frame. Only the roles for the three main dream characters, Betty, Rita, and Adam, are shown. Certain frames are introduced without explication, assuming that the reader knows the primary roles and scripts of the frame.

(3)
SMALL-TOWN GIRL > *MOVIE PRODUCTION* > *CORRUPTION* > *MAFIA* > *EXTORTION/VICTIM* >
ADAM

This is one of the Adam plots: the Mafia force Adam to cast d-Camilla Rhodes as the female lead in his film. This serves goal 7, revenge against Adam. Table 10 lays out the mappings in this plot thread of the *MAFIA* frame. Since *MAFIA* does not appear in either Input, only inner-space VRs are shown.

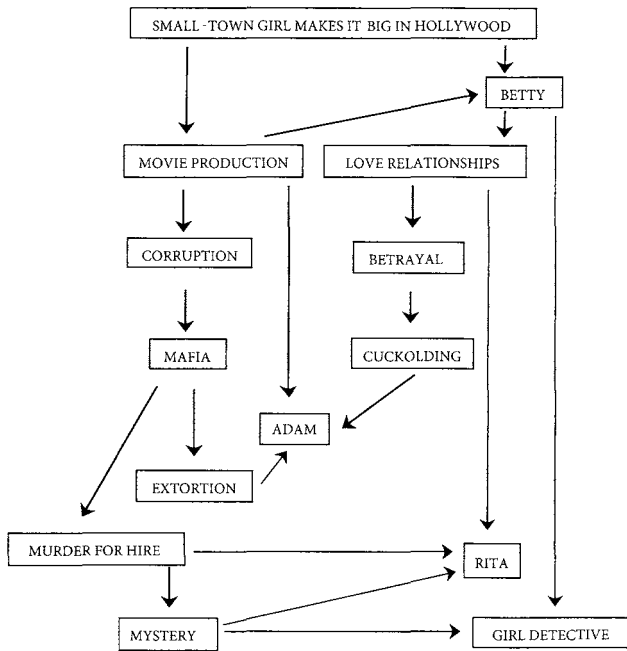


Figure 3. Hierarchical relations among frames/elements

It is worth noting that the female lead is the only aspect of the film that Roque wants to control. The Cowboy tells Adam, “The rest of the cast can stay. That is up to you. But that lead girl is not up to you.” r-Diane’s grudge against Adam has to do with his choice of leading lady for *The Sylvia North Story*, so this power is what is taken from him in the dream.

Another reason for importing *MAFIA* into the dream is to corner Adam: if you do not comply with the Mafia, or if you go to the police, you may lose everything, even your life. r-Diane chose the most powerful avenger she could, short of an authoritarian dictator or God. She Changes the scale of power by magnifying the intensity of the power asymmetry (between Adam and her) of Input 1. As is often the case in such circumstances, r-Diane’s abject powerlessness remedies itself in fantasies of a far more powerful avenger.

r-Diane suffered not only cuckolding and rejection at the hands of Adam, but also humiliation. These tables must also be turned, and they are, increasing the intensity of

Table 10. Mappings in the *MAFIA/EXTORTION* frame

Roles/Outcomes	VR	Blend Dream
Perpetrator	<i>Role/Value</i>	Roque
Subcontractors?	<i>Role/Value</i>	Mafiosi; Cowboy
Target	<i>Role/Value</i>	Adam
Motive	<i>Variable binding</i>	Control lead role
Means	<i>Variable binding</i>	Threats: bankruptcy; career destruction; personal harm/death
Outcome	<i>Variable binding</i>	Acquiescence

the humiliation by adding absurd details: these Mafiosi in the Blend are outrageous to the point of slapstick. One snarls over his briefcase and rants in a high temper when Adam rejects “the girl”. The other, who apparently has a reputation for being extremely hard to please, spits out the espresso he is served like a cat spits out an unpleasant morsel. The displeasure of these Mafiosi strikes terror into the film agents as well as the server: they quake and apologize abjectly.

Then there is the Cowboy, a major element of Lynch's use of the hermeneutic code. The Cowboy is known only by this epithet, and is clearly a major power in the EXTORTION scheme. Adam is ordered to meet him in a corral at the end of a canyon road late at night, on the day of his disastrous encounter with the Mafiosi.

Oblivious to his own odd looks – pale and pink as a baby, dwarfed by his ten-gallon hat, and wearing an unfashionably retro coat that looks to have been made out of old camp blankets – he rides a very high horse indeed. In a calm, quiet voice, he puts Adam in his place: “You're too busy being a smart aleck to be thinkin'. Now I want ya to think and stop bein' a smart aleck. Can you try that for me?” Then, in his homespun accent, he speaks in didactic metaphors: “There's sometimes a buggy. How many drivers does a buggy have? ... So let's just say I'm drivin' this buggy ... and if you fix your attitude, you can ride along with me.”

Hipster Adam finds all of this ridiculous, but is cowed by the blunt power of these men. Later, in the audition scene, he slumps sullenly in his director's chair, watching the dead-eyed d-Camilla sing a '50s rock tune; he utters, “This is the girl.”

4.4 Murder for hire

Above, I posed the question of why d-Diane's death and MURDER FOR HIRE remain in the dream. r-Diane is horrified that she was capable of having her beloved killed; guilt oppresses her. She fears getting caught (detectives want to question her). And, in spite of her rage towards r-Camilla, she still loves her, and grieves her death. These emotions, plus her humiliating acting failure, are driving her towards psychosis. Her dream must provide relief from these torments.

I noted in Section 3 that Betty frees r-Diane of responsibility for r-Camilla's death. To preserve the disanalogy between Input 1 and the Blend down to the last detail, d-Diane, the corpse on the bed, should be the instigator of the hit Rita escapes. A connection is suggested by the fact that the chain of phone calls resulting from Camilla/Rita's narrow escape ends at Diane's phone. Other facts make her seem peripheral, however. The circumstances of the attempted murder indicate considerably greater resources than someone of r-Diane's caliber could muster: the well-dressed hit men who drive the limo on Mulholland Drive; the grand-looking interior from which Roque is telephoned; Roque's own wealth; the G-man-style suits looking for Rita in a large sedan. All of this belies considerably more sophistication than one young woman paying a shabby hit man a few thousand dollars to send r-Camilla to her doom. Table 11 shows the dream's revision of MURDER FOR HIRE. Input 2 is not shown in this table, because r-Diane wishes that the hit had never been planned.

Table 11. Mappings in the MURDER FOR HIRE frame

Roles/Outcomes	Input 1 Reality	VR	Blend Dream
Instigator	D	<i>Ds</i>	Roque
Target	C	<i>Id</i>	C
Motive	Rage over rejection	<i>Ds</i>	Witness elimination?
Killer(s)	Joe the hit man	<i>Ds</i>	Limo driver
		<i>Id</i>	Joe the hit man
Outcome	C dies	<i>Ds</i>	C escapes miraculously

Once again, the schematic variables map to the Blend, but crucial Role Values differ. Importantly, Intentionality is once again Disanalogously projected onto a different instigator; not even d-Diane bears full responsibility for the murder plot.

If d-Diane were the sole culprit in the murder plot, why should Roque care that “the girl is still missing”, almost certainly a reference to Rita? How would her death serve Roque’s interests? We can only speculate. In reality, Camilla was Adam’s choice for the lead in *The Sylvia North Story*; apparently, Diane auditioned so poorly that she did not have a chance at the role. As part of Diane’s revenge, this choice must be taken away from him: both the real Camilla and control over her replacement. Perhaps Roque paid Camilla/Rita off to withdraw from the audition and keep quiet (hence the wads of cash in Rita’s purse), intending all along to have her killed and recover the money. Perhaps her fateful limo ride was ferrying her away from a meeting with the Mafia to which she had been instructed to come without any identification. Transferring guilt for the murder plot to the Mafia also helps dilute the degree of evil such a plot implies for an “ordinary” person like Diane Selwyn. The Mafia are known for their ruthlessness and total lack of conscience; such murders are virtual tradition for them. The simultaneous ruthless premeditation and cowardice behind a hired murder adds significant seriousness to a one-on-one “civilian” murder, making it less mitigable than an unpremeditated, spur-of-the-moment crime of passion.

Motive for the murder notwithstanding, the essential purpose of these disanalogies is to relieve r-Diane of her guilt for r-Camilla’s murder. The disanalogous outcome between reality and the Blend undoes her grief, and permits Camilla, as Rita, to start anew and find love in the arms of Betty Elms.

(4)

SMALL-TOWN GIRL > MOVIE PRODUCTION > CORRUPTION > MAFIA > MURDER FOR HIRE/
VICTIM > RITA

This is the thread of the Mafia plot which provides alternate culprits for the hit on r-Camilla. This plot leaves a hermeneutic thread hanging: why do the Mafiosi want r-Camilla dead? I proposed above that they paid r-Camilla a large sum of money (the wads of cash in her purse) to bow out of auditions for Adam’s film so that they could (for reasons also left unexplained) place another actress in the leading role.

The MURDER FOR HIRE thread enables MYSTERY via another popular dramatic frame: NARROW ESCAPE. In this case, NARROW ESCAPE Uniquely includes the AMNESIA frame, Caused by the car accident, which makes Rita’s identity and story a MYSTERY.

(5)

... MURDER FOR HIRE > NARROW ESCAPE > AMNESIA > MYSTERY

This plotline makes Camilla/Rita connectable to Betty by binding her to a Role in both the LOVE RELATIONSHIPS and GIRL DETECTIVE frames. The goals served here, as noted above, are 2 (requited love), 3 (undo murder, hence grief), 4 (d-Camilla gets a chance to reform herself), 7 (revenge against d-Camilla: the Mafia danger keeps Rita frightened, a punishment).

(6)

... MURDER FOR HIRE > NARROW ESCAPE > AMNESIA > MYSTERY > GIRL DETECTIVE/
BETTY

The GIRL DETECTIVE Role enables one aspect of Betty's adventurousness, cleverness, and charisma, which not only makes her attractive, but allows her to be dominant in her relationship with Rita. This serves goal 6 (reversal of the interpersonal power force dynamics between r-Diane and r-Camilla).

(7)

LOVE RELATIONSHIPS > LOVERS/BETTY, RITA

Table 7 and previous sections have laid out how the dream rearranges love relationships by reincarnating r-Diane as Betty, and having Camilla both escape the hit and change such that she is vulnerable to Betty's charisma.

(8)

LOVE RELATIONSHIPS > BETRAYAL

Also shown in Table 7 are the Unique happenings between Adam Kesher and his wife, Lorraine. This thread is made up of whole cloth just to preserve the topology of betrayal, but have it turned on Adam. Lorraine's affair with the pool man compounds Adam's humiliation at the hands of the Mafia. This subplot is r-Diane's personal class war. Her experience in Hollywood included being humiliated by the wealthy, powerful insiders; they cared nothing for the dreams of a small-town girl. Adam's directorial success, luxury home, trendy glasses and clothing can't compete with the tattooed working-class hunk (Billy Ray Cyrus). Adam's reaction is childish – he pours pink paint on Lorraine's jewelry. He then loses a fight with Gene and is thrown out of his own house, taking refuge in a flea-bag hotel far shabbier than r-Diane's humble apartment. How the mighty are fallen.

4. Conclusion

Mulholland Drive is a CB tour-de-force. The interweaving of frames in the Blend, with the added structural core of MAFIA, produces complex interlinkings of Cause and Effect that neatly and coherently achieve r-Diane's psychological goals. But Lynch disguises all of this in order to engage the viewer deeply in experiences projected from his protagonist to his audience.

This paper has treated only some major aspects of the film, in hopes of revealing its true coherence. I aimed, as well, to demonstrate the value of CB theory as a tool for revealing meanings presented through modern cinematic techniques such as Lynch's. Some might say that such dissection destroys the mystery of the film, but, for me, it reveals the consummate artistry and attention to detail of Lynch's particular genius. The paper also shows that CB is as excellent a stylistics tool for film as it is for literature.

Lastly, the paper suggests application of CB to psychological theories of dreaming. In further work, I would like to explore the similarities between *Vital Relations* and the defense mechanisms of psychoanalysis. Diane's dream involves feats of projection, splitting, idealization, and reaction formation, all analogs of various sorts of disanalogy, scalar manipulation, compression and decompression, and so on. If CB truly is a "theory of everything (cognitive)", such explorations should prove fruitful.

References

- Bulkeley, K. 2003. Dreaming and the Cinema of David Lynch. *Dreaming: The Journal of the Association for the Study of Dreams* 13 (1): 49–60.
- Cross, D. n.d. *Public Servant Lifestyle*. http://www.publicservantlifestyle.co.uk/dynamic/sections/entertainment/article_display.php?id=2368
- Domhoff, G. W. 2001. Why did empirical dream researchers reject Freud? A critique of historical claims by Mark Solms. *Dreaming* 14: 3–17.
- Fauconnier, G. & Turner, M. 2002. *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. New York: Basic Books.
- Felluga, D. 2003. Terms Used by Narratology and Film Theory. *Introductory Guide to Critical Theory*. <http://www.cla.purdue.edu/academic/engl/theory/narratology/terms>
- Kantrowitz, B. & Springen, K. 2006. What dreams are made of. *Newsweek*, August 9 2006.
- Lakoff, G. 1993. How metaphor structures dreams: The theory of conceptual metaphor applied to dream analysis. *Dreaming* 5 (2): 77–98.
- Lakoff, G. 1997. How unconscious metaphorical thought shapes dreams. In *Cognitive science and the unconscious*, D. Stein (ed.), 89–120. Washington, DC: American Psychiatric Press.
- Rodley, C. (ed.) 2005. *Lynch on Lynch*. London: Faber and Faber.
- Semino, E. & Swindlehurst, K. 1996. Metaphor and mind style in Ken Kesey's *One Flew Over the Cuckoo's Nest*. *Style* 30: 143–166.
- Steen, F. F. 1998. 'The Time of Unrememberable Being': Wordsworth's autobiography of the imagination. *Auto/Biography Studies* 13 (1): 7–38.
- Turner, M. 1996. *The Literary Mind: The Origins of Thought and Language*. Oxford: Oxford University Press.

“I was in that room!”

Conceptual integration of content and context in a writer’s vs. a prosecutor’s description of a murder*

Esther Pascual

1. Introduction

Cognitive Linguistics has shown time and again that fictivity and imagination play an important role in human language. In fact, counter to the folk belief, the work of cognitive linguists suggests that fictivity and imagination play as big a role in ordinary as in literary language. The basic claim is that metaphors (cf. Lakoff and Johnson 1980, 1999; Lakoff and Turner 1989) or non-factual scenarios (cf. Fauconnier 1994, 1997) for instance are not solely devices of the poetic imagination and the rhetorical flourish. They reflect the way our mind work and their use is thus not restricted to literature.

In this paper I assume this view in the comparison of the draft of a non-fiction book on a high-profile murder case with the prosecutor’s closing argument at the trial for the same murder. The focus of study is on two particularly dramatic fragments of the novelist’s book and the prosecutor’s argument. These deal with the killing of the victim. I have chosen the description of this particular event basically because no aspect related to the murder had been irrefutably proven. As I hope will become evident, as a result of the lack of conclusive tangible direct evidence in this case, both the writer’s description and that of the prosecutor set up conceptual configurations that were halfway between fact and fiction. More specifically, I discuss the description of the assumed criminal facts through a conceptual integration of the situated context of ongoing communication with the discourse content.

* The research presented in this paper was supported by post-doctoral Veni fellowship 275-70-012 from the Netherlands Organization for Scientific Research (NWO). The ethnographic data collection was funded by a Fulbright grant (M081000, 2000–2002). I am particularly grateful to the subjects of this study, novelist A. W. and Deputy District Attorney D. G., who generously agreed to be interviewed and allowed me to use and publish the data collected. Special thanks go to A. W. for being willing to share with me an electronic version of the first draft of her non-fiction book for analysis and to KSNB television for trusting me with a raw videotape of the prosecutor’s argument. I am most grateful to Theo Janssen, Kevin Moore, Kashmiri Stec, Eve Sweetser and Arie Verhagen for useful comments and suggestions for improvement. All shortcomings are of course my own. Contact details: Department of Communication and Information Sciences, Faculty of Arts, Groningen University, Oude Kijk in ‘t Jatstraat, 26, 9712 EK Groningen, The Netherlands. Tel.: +31 (0)50 363 8133, Fax: +31 (0)50 363 6855.

This study is framed within Fauconnier's (1985) *Mental Space Theory*, subsequently elaborated into *Conceptual Integration Theory*, also known as 'blending' (Fauconnier and Turner 1998, 2002). The theory of Mental Spaces and Conceptual Blending was first mainly used to explain processes of meaning construction in isolated sentences or short paragraphs. Following the latest developments of the theory (cf. Oakley 1998; Coulson and Oakley 2006; Dancygier 2006; Hougaard and Oakley 2008), I intend to show that it can also account for situated pieces of oral discourses and written texts.

2. Mental Spaces and conceptual blending

Mental spaces (Fauconnier 1985) constitute abstract mental constructs or conceptual domains that are dynamically prompted as discourse unfolds. Examples of these are: the world defined by a picture, a world of fiction, the world of a person's beliefs and desires, hypotheticals, or time spans. Fauconnier's theory can elegantly solve intricate problems of ambiguous linguistic reference, by pointing at the possible conceptual mappings between an element in one mental space and its counterpart(s) in another space. Fauconnier and Turner (1998, 2002) have further shown that elements in different mental spaces can become fused into one single element with properties from the so-called 'input spaces' and emergent properties of its own. Consider for instance an advertisement run by an education partnership, discussed in Fauconnier and Turner (2000, 2002). In this ad, three children dressed as surgeons appear in front of a patient lying in an operating room. The headline reads:

- (1) Joey, Katie and Todd will be performing your bypass.

Two time spaces are involved in the configuration set up by this ad, a Present and a Future space. In the Present space, Joey, Katie and Todd are about seven years old and have no competence as surgeons. In the Future space, Joey, Katie and Todd are operating on the reader. Of course, the ad does not suggest that at a given time, say three days after reading the ad, the reader will be operated by these children. Rather, what needs to be interpreted is that in the Future space, after having received their education, the children in the Present space will be adult professionals, who will have to perform difficult tasks requiring great competence and skill. This becomes clear once we consider the body of the ad, which reads:

- (2) Before you know it, these kids will be doctors, nurses and medical technicians, possibly yours. They'll need an excellent grasp of laser technology, advanced computing and molecular genetics. Unfortunately, very few American children are being prepared to master such sophisticated subjects. If we want children who can handle tomorrow's good jobs, more kids need to take more challenging academic courses. To find out how you can help the effort to raise standards in America's schools, please call 1-800-96-PROMISE. If we make changes now, we can prevent a lot of pain later on.

In "Joey, Katie and Todd will be performing your bypass", a conceptual blending of the Present and the Future spaces occurs, in which the children in the Present space are

mapped onto and blended with their own selves in a (potential) Future space. In the ad's picture, they look like the children they are today, but are engaged in the tasks they may be in charge of tomorrow. Note that yet another blend is involved here. By reading the ad, the reader becomes part of the message, as s/he is conceptually integrated with the patient to be operated on by the three children. This identity blend is a critical part of the ad's meaning and persuasive function. It is after all by imagining oneself being treated by unskilled professionals, that one can fully realize the importance of investing in education. Such an integration of an individual in the communicative context and a character in the message content is not uncommon in print advertisements and television commercials. (For a blending analysis of another example, see Coulson and Pascual 2006: 157–159.) More generally, the content-context blend is particularly common in different sorts of persuasive discourse, as well as in types of discourse in which the author aims at the addressees' emotional involvement in the story.

In this paper I examine fragments of emotionally charged discourse which prompt a conceptual integration of content and context. This type of blend invites the analyst to look further into the role of the context of production and interpretation, i.e. the Here-and-Now space,¹ in the overall configuration of the Current Discourse Space (Langacker 2001), namely the conceptual domain comprising "those elements and relations construed as being shared by the speaker and hearer as a basis for communication" (2001: 144). I believe that the study of the Current Discourse Space should go beyond the configuration set up by the verbal register of a communicative act or what I call the Verbal space (Pascual 2002: 82). As it happens, cognitive linguists in general and conceptual blending scholars in particular have only recently started to include the overall situation of communication as part of their analyses. In his study of American Sign Language, for instance, Scott Liddell (1995, 1998) introduces what he calls the 'Real space'. This is the mental space of the conceptualizer's perceived physical surroundings. Liddell shows that our perception of the immediate environment can be an input to a conceptual integration network involving other spaces, producing a so-called 'grounded blend'. In this space, a gesturer's arm and hand, for example, can function as a so-called 'surrogate', as they are understood as standing for Garfield the Cat during the narration of a cartoon. Similarly, Chris Sinha (2005) analyzes the complex blends involved in young children's symbolic play, which show "the socially collaborative, culturally and materially grounded nature of the human mind" (ibid.: 1537). Along very much the same lines, cognitive anthropologist Ed Hutchins (2005) has introduced the notion of 'material anchor', namely "an input space from which material structure is projected" (Hutchins 2005: 1555). This space is set up by a cognitive artifact or piece of material structure serving to stabilize the representation of conceptual relations. Examples are maps, calendars or one's own fingers used in counting. In this paper I use the notions of Real space, surrogate, and material anchor in the analysis of the conceptual networks involving the integration of the Verbal space with the

1. The Here-and-Now Space corresponds to what Brandt and Brandt (2005) call the 'Semiotic Space'. I have decided not to use this term, finding it not very accurate to characterize the situation of ongoing communication.

Here-and-Now space. Following Williams (2004), the grounding character of a space is represented in my figures with a square behind the space in question.

3. Data

This paper is based on ethnographic data from a high-profile murder trial that I observed in a Californian county court in 2000 (Pascual 2002). The defendant in this case was a financial manager accused of brutally killing his wife in the couple's bedroom, so he could collect her pension and three life-insurance policies, of which he was the only beneficiary. No clear evidence or alibi was provided to prove the defendant not guilty and he was the only witness for the defense. Despite the transparent incriminatory nature of the evidence – admitted by the chief deputy defense attorney in my interview with him (Int.9-DC: 10–11) – the prosecution team could find no direct evidence against the defendant. After a two-week trial the defendant was found guilty and later sentenced to life in prison without parole plus seven years.

Analysis is informed by a draft of the non-fiction book on the case; a thirty-minute videotape of the prosecutor's argument at trial borrowed from a television station; and my feedback interviews with the amateur writer who wrote the manuscript (Int.8-Nov) and the prosecutor whose closing argument is discussed (Int.1-DA). Also, extensive ethnography was carried out, mainly involving direct observation of the trial and sentencing and in-depth interviews with four trial participants and four attendees, including an alternate juror. For privacy reasons, only initials are used. In the examples to be discussed, direct quotes from the prosecutor's speech come from the official court transcript, enriched with minimal paralinguistic information (e.g. [points to his left]) and clarifications (e.g. [victim]) in square brackets, and italics for prosodic emphasis (e.g. '*why?*') from my ethnographic notes and from the videotape. In both the example from the prosecutor's speech and the writer's work relevant words and expressions are underlined.

4. Analysis

This section applies the theory of Mental Spaces and Conceptual Blending to the analysis of the writer's narration of the crime in the draft of her non-fiction book and the prosecutor's description of the same event to the jury at trial. These were most probably prepared independently from each other, since the novelist wrote the relevant chapter before the trial started and it is highly unlikely that the prosecutor ever read the writer's manuscript.

The two relevant fragments dealing with the victim's death are interesting, since the assumptions they were based on had not been irrefutably proven. The crime occurred in the privacy of the couple's bedroom, with no eyewitnesses who could explain what happened (apart from the defendant, who did not admit guilt, and the victim, who did not survive the attack). Needless to say, there were no audio- or video-recordings of the criminal events either. As a matter of fact, the defendant first stated that the victim had died as a result of a medical accident, and hence that no murder had occurred whatsoever.

The two pieces of discourse to be discussed are also interesting for comparison purposes, since – while belonging to different genres – they both deal with the same objective event in the real world and are roughly based on the same information. (It should be noted that the writer, an amateur novelist familiar with the case, followed the entire trial and media coverage, interviewed the victim’s family and even visited the scene of the crime.) Moreover, even though the ultimate communicative goal of books and legal arguments are radically different, in this case both the novelist and the prosecutor had as one of their main objectives the achievement of sympathy for the victim. The novelist wrote her book as a tribute to the woman’s memory (Int.8-Nov:5) and wanted the reader to feel what she had gone through (Int.8-Nov:8). The prosecutor understood the trial as “a battle for sympathy”, which he could only win if the jury empathized with the victim (Pascual 2002: 116, 2008). I suggest that in both cases sympathy is achieved through a conceptual integration of the Verbal space and with Here-and-Now space.

4.1 Setting the conceptual scene

In a four-hour-long interview, the amateur writer who wrote the non-fiction book on the case at hand told me that the lack of direct tangible evidence on the circumstances surrounding the victim’s death made her wonder about the truth of every aspect surrounding the attack (Int.8-Nov:9). She also said that this had led her to give her manuscript the title “*It Remains To Be Seen*” (henceforth *IRTBS*). Similarly, the title of the chapter in which the murder is narrated does not appear in the form of an assertion, but a question: “What Ever Happened to R. C.?” Not too surprisingly then, the entire description of the murder is first introduced through setting up a non-committal Imaginary rather than a Reality space as follows (*IRTBS*, ch. 3: 11):

- (3) Come. Let us use our imaginations, shall we? Soar back in time to the balmy summer’s eve of August 13th, 1999. Pretend we are flies on the proverbial wall, or innocents, perhaps, armed with a telescope for a little night gazing...

This piece of text serves as a space-builder. It creates an imaginary scenario in which the events to be narrated are a plain Past Reality space. Through the KNOWING IS SEEING metaphor (Lakoff and Johnson 1980), the rhetorical metamorphosis prompted turns writer and reader into direct ‘observers’ of the crime, as it may have occurred. This construction allows the narration of the events as though they were unfolding in front of the writer and reader. The storytelling is conceptually integrated with the story being told. In other words, the narrating in the Here-and-Now space is presented as occurring in the same mental space as that which is narrated, which inhabits the Verbal space. This image emerges from the conventional blend in which reading is conceptualized as the writer speaking to the reader directly (Herman 1999; Fauconnier and Turner 2002: 210–211). In the narrative, the criminal facts are recounted as though writer and reader were ‘perceiving’ them through their mind’s eye (see Figure 1).

This conceptual configuration succeeds in turning hypothetical past events into present-time ones. The story can then be presented vividly with the use of exclamatives, deictic

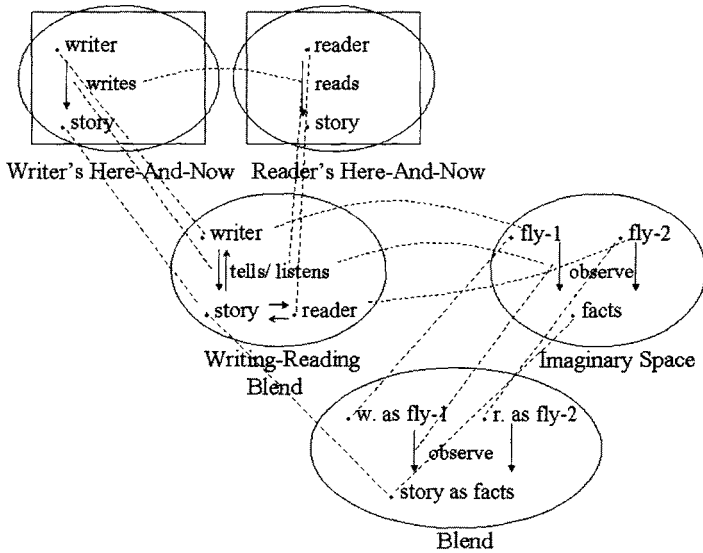


Figure 1. The writer and reader as imaginary eye-witnesses

pronouns, and verbs of perception such as ‘see’, ‘watch’, ‘listen’, and ‘feel’ in the imperative mode. Consider for instance (*IRTBS*, ch. 3: 11–14):

- (4) You can see it, no? [...] feel it! There! Now! A sinister fillip rippling tranquility [...] Her bloodied head moves, twitches [...] R. is still alive! Watch her chest rise and fall. And listen! [...] And, yes, we creep forward then, invisible in the searing hatred, for a better view. My god!

Through the use of the present tense, which seems to function as a historical present, the events appear narrated as in a novel. Note too that when working on her manuscript, the writer could only have had a ‘virtual reader’ (Langacker 1999: 95) in mind, as opposed to any individual one. This notwithstanding, in the Imaginary space she set up, the writer finds herself in the crime scene together with the actual reader who happens to be going through her manuscript in a given space and at a given time. This allows her to address that reader directly with the second person pronoun and imperative forms. By so doing, the reader is not treated as an overhearer, as is often the case in published texts (Tobin 2006). The reader becomes an addressee as well as an onlooker of the events.²

2. As a fictive onlooker, the reader of the non-fiction novel does not become part of the story being read, which would be rather peculiar, since the narrated events are based on real life. Therefore, even though the content-context blend involved has an important emotional function, it is not essential for understanding the main plot. Some advertisements, such as the bypass one, as well as cartoons, short stories and novels are fundamentally based on a content-context blend. An example is Julio Cortázar’s short story “Continuity of Parks”, in which the protagonist, a reader immersed in a thriller, is first “witness” to a narrated encounter in a mountain cabin and subsequently becomes the victim of the murder mystery he is reading.

Let us consider now how the prosecutor introduced his description of the criminal facts to the jury in the trial for the same murder. Bearing in mind that the prosecutor has the ultimate proof of guilt, his main objective is not only to tell a tragic story, as in the writer's case, but to have the defendant found guilty of the charges. Thus, he cannot be satisfied by merely presenting the crime within a Hypothetical space – let alone an Imaginary one – as the writer did. Even if in actuality the prosecutor also operates largely on the hypothetical, he needs to account for each and every one of the elements and events in his Hypothetical space. Not surprisingly then, the prosecutor's conceptual framing of the crime in his discourse to the jury was less imaginative, but certainly not less complex than the writer's. He first warned the jury that (Vol. 6, 1355:23–28):

- (5) until you can see that killing, until you can see what the defendant did to her and contemplate it and understand it, you cannot judge the defendant's conduct. Can you go back to [...] that night, and see R. C. being struck in the head, struck in the head while she lay in bed?

In this fragment the prosecutor is asking the jury to observe the criminal events (for which there was no direct evidence) through their mind's eye, using the same verb of perception 'to see' as the writer did in (4). The knowledge of the events and how they must have occurred should be such that, through the KNOWING IS SEEING metaphor (Lakoff and Johnson 1980), the jury may be able to watch them happen. He subsequently started his description of the attack as follows (Vol. 6, 1356:9–10):

- (6) The master bedroom and the master bathroom [i.e. crime scene] are the size of this courtroom [extends arms widely] together.

With this simple utterance, the prosecutor establishes an analogical relation between the courtroom in the Real space, that is the physical space where the ongoing discourse is being delivered, and the crime scene, where the narrated events in the Verbal space took place. This integration is motivated by the immediacy of the prosecutor's discourse, since its communicative situation is one which – counter to the writing and subsequent reading – is shared between addresser and addressees. The establishment of such a relation between communicative context and content prepared the discourse ground for the deictic use of spatial pronouns in the following fragment (Vol. 6, 1356:16-1358:4):

- (7) You can tell from what B. K. [blood spatter expert] said that the original attack probably occurred around there [points to his left] because you don't see so much bloodletting at that point. [...] She ends up landing in an area approximately right here [stands in middle of courtroom] in the middle of the master bedroom, and she is still being hit.

The pronouns "there" and "here" refer to locations of the event recounted. However, the speaker's pointing and moving around suggests that these pronouns are not used anaphorically, but deictically. Once a size mapping has been established between the situated context of the discourse production and an element in the story being told, the prosecutor can use the one, i.e. the immediate physical space of the courtroom, to refer to spatial relationships that hold in the other, i.e. the bedroom in which the narrated murder took place. The courtroom in the Real space has therefore become a 'material anchor' for a content-context

blend. Furthermore, just as was the case in the writer's manuscript, the use of the present tense in (7) helps to construe the events as in a narration. The shift to the past simple indicates a shift to the Here-and-Now space, from which facts in the Past Real space are recounted. The events told in the present progressive are construed in the blend of what Dudis (2004b) calls 'depicting time'.³ The prosecutor also acts as what Liddell (1995) calls a 'surrogate' of the victim, moving around the courtroom as she must have moved around the crime scene. The time progression of the criminal facts depicted is therefore mapped onto Real time, that is, the time of discourse production in the Here-and-Now space. Consider now the following piece of discourse (Vol. 6, 1357: 2–5):

- (8) And she gets into the corner. [walks to corner of courtroom] She's got to get into this corner at D4 [points at picture exhibit].

Here, the prosecutor is not only conceptually integrating the physical space he and his addressees and overhearers find themselves in (i.e. the courtroom) with the physical space within which his story occurred (i.e. the crime scene). He is also integrating in the overall configuration the Picture space of a crime scene photograph ("D4"), which was taken at a different time, after the occurrence of the narrated facts. This Picture space, which is conceptually linked to the Post-Crime Past space of the investigation, is being accessed by the jury when the prosecutor shows it to them in the Here-and-Now space. Therefore, the exhibit serves as a 'material anchor,' perceived by the jury in their Real space. This overall configuration allows the prosecutor to refer to actual locations in the crime scene (now conceptualized as the courtroom) through reference to pictures of these locations taken by the investigators.

Interestingly, once he has taken himself and the results of the investigation into the picture, the prosecutor also gets the defendant into that conceptual network. Reaching the end of his argument, the prosecutor adds (Vol. 6, 1357: 27–28; 1358: 11–12):

- (9) She [victim] pushes over a chair, or the defendant does, and she's still hit in the same area. [...] That man over there [points at defendant] had incapacitated her at D3.

In this apparently simple piece of discourse the prosecutor is setting up a complex conceptual configuration involving: i) an individual in the Real space of the courtroom (i.e. the defendant); ii) an individual in the Past space of the crime (i.e. the murder victim); and iii) an entity in a Picture space (i.e. the exact location where the murder was finalized). Thus, the prosecutor is implicitly setting up a Trial-Investigation-Facts blend, which presupposed an identity relation between the person accused of the crime and the person responsible for it. The overall cognitive configuration set up by this piece of discourse is represented schematically in Figure 2.

Through setting up this multimodal grounded blend, the prosecutor succeeded in re-enacting the crime before his addressees, in an even more complex manner than the writer

3. Dudis (2004a, b) has shown that Real space blends demonstrating actions are indispensable to American Sign Language discourse. Just as the prosecutor did in his closing argument, signers integrate conceptualized event time and discourse time for purposes of iconically depicting events as if they were occurring in the Here-and-Now.

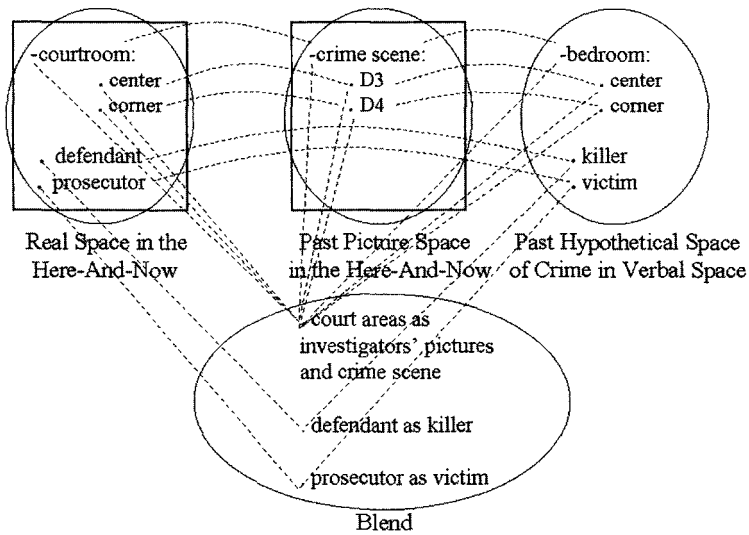


Figure 2. The courtroom as the scene of the crime

did in her non-fiction book. Just as the novelist’s use of deictic forms, the prosecutor’s illustrates the dynamic character of the indexical ground (see also work by Bühler 1934 and Hanks 1990, 1992). The prosecutor’s discourse also shows that the discourse building process may not only be facilitated by gestural information, as has often been demonstrated (cf. McNeill 2000, 2005; Enfield 2003; Luchjenbroers 2006), but also by proxemics. When I asked a juror how effective she thought the prosecutor’s discourse was, she exclaimed: “he put us there, I was in that room!” (Int.7-Jur: 13).

4.2 Fatal strikes with the missing weapon

At the beginning of the previous section a fragment was discussed in which the writer invites the reader to join her in a mental journey back in time to the evening of the murder. It was suggested that this involved the conceptual integration of the writing and reading events with the story being written or read into a scenario in which writer and reader observe the events in question as they develop before them. In the blend these events are directly perceived, hence becoming concrete and real. Note, however, that since writer and reader are only to “pretend” to be able to observe the crime directly, an act created by the use of their “imagination”, the writer’s account to follow cannot have the same degree of reliability as that of an actual eyewitness in the Past Reality space of the crime. In fact, given that the ‘observation’ of the crime occurs within an Imaginary space, the crime itself, that is, what is observed through the mind’s eye, may be conceptualized within a Hypothetical – and not necessarily a Past Reality – space. The exact correspondences between this Hypothetical space and the Past Reality space of the actual facts need not be specified. This particular conceptual configuration allows the writer to

present details surrounding the fatal events even when there was no absolute certainty that these details – and in fact the entire event narrated – actually took place. Consider first (10) below (*IRTBS*, ch. 3: 12, 13):

- (10) With a violence shattering this placid connubial scene forever, the weapon whips, [...] the terrible whistle as the murder weapon cleaves the air into jagged shards. [...] The wicked weapon strikes the bed with a crisp whomp [...] Blood spins through the air now, flowing down the length of the weapon [...] Without a beat of hesitation the weapon streaks through the air, its edge striking the woman's head as she lies on the floor.

In this fragment a murder weapon appears in the scene, interacting with the furniture and the victim's blood and body. In actuality, no weapon was ever found by investigators which could be related to the crime. However, investigators did infer from the victim's wounds and the blood spatter on the bed and walls that a weapon must have been used. Thus, in (10), the weapon in the investigators' (Professional) Belief or Inference space – which was presented in the Past Reality space of witness testimony – appears through its counterpart in the Hypothetical space of the facts. More specifically, since the traces left by what must have been the weapon indicated that it must have been a long metal object, investigators testified that they believed the weapon was most probably a fireplace poker. Critically, the couple owned a poker, which went missing right after the crime. Since one would expect the perpetrator of the crime to try and hide the evidence for it, the weapon was strongly suspected to correspond to that missing poker. Thus, the overall underlying conceptual configuration upon which (10) is based involves a missing-X construction blend (Fauconnier and Turner 2002: 241ff.). The weapon inherits thing-hood from the Past Pre-Crime space in which there is a poker by the fireplace. This space was presented in the Real space of ongoing communication through showing the jury an old picture of the fireplace area with the poker hanging on the wall. The missing poker inherits its physical characteristics of being a gap from the 'actual' input, in which there is no poker in the home. This missing poker element in the Blended space is thus a compression of the disanalogy between the Post-Crime Reality space of the trial and the Pre-Crime Reality space (see Figure 3). Once compressed, speakers can refer to the missing poker and map it onto the weapon that must have been used in the Past Crime space. Again, in the writer's narrative, the weapon in the investigator's (Professional) Belief or Inference space is implicitly mapped onto the weapon in the book's Imaginary space. Since that element is not conceptualized within a Reality space it need not be introduced and referred to in the text with a long non-committal description such as "what was probably a poker." Take the following fragment (*IRTBS*, ch. 3: 13):

- (11) ...the fiend leaps onto the raised hearth of the fireplace [...], and with calculated malice backhands her across the face with the poker. [...] The poker connects with her head; tissue, hair and blood clot the air, fly up through the open transom at the top of the shower stall. Pound! The poker strikes R.'s face, directly across her eyes and nose, splintering her nasal bones.

Consider now the way in which the weapon is presented by the prosecutor in his discourse to the jury in the case at hand (Vol. 6, 1363: 11–18):

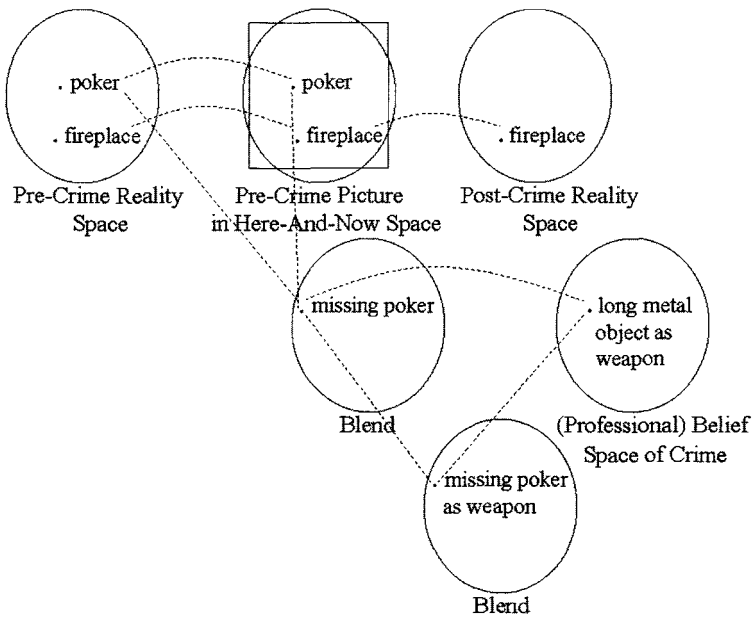


Figure 3. The missing poker as the murder weapon

- (12) On People’s 11, 11A, at the edge of this bedspread, that’s the weapon right there. There’s a linear object. It’s approximately three feet in length and it lies across the bedspread. It is covered in blood. But that linear object cannot be explained by any other device at the crime scene except it is consistent with a fireplace poker and the fireplace tools are gone. That explains that mark.

In this extract the prosecutor discusses one of the pictures of the crime scene taken by investigators. Hence, what is characterized as “the weapon” in “People’s 11A” is in fact a picture representation from the Post-Crime Past Reality space of the investigation, perceived by the jury in their Real space. Significantly, this is not a picture of an “actual linear object”. Rather, it is the photograph of a long mark of blood on the couple’s bed. Thus, when the prosecutor says “that’s the weapon right there”, the connection between the actual weapon used at the time of facts, possibly a fireplace poker, and the picture taken during the investigation is not merely an analogical mapping between representation and thing represented. Since the weapon was never found, a picture of it could not be taken by investigators. Rather, there is a connection chain from the real blood mark to its representation in the picture, and then from this representation to the actual object that left that mark. An EFFECT FOR CAUSE metonymy (Panther and Thornburg 2000) is therefore involved, since the blood mark stands for the long object stained with the victim’s blood, which for this very reason could only be the murder weapon. At the same time, since the mark is consistent with the mark that a fireplace poker could have left, and since the couple’s fireplace tools went missing after the crime, an identity relation is set up between the weapon in the prosecutor’s – and the investigator’s – (Professional)

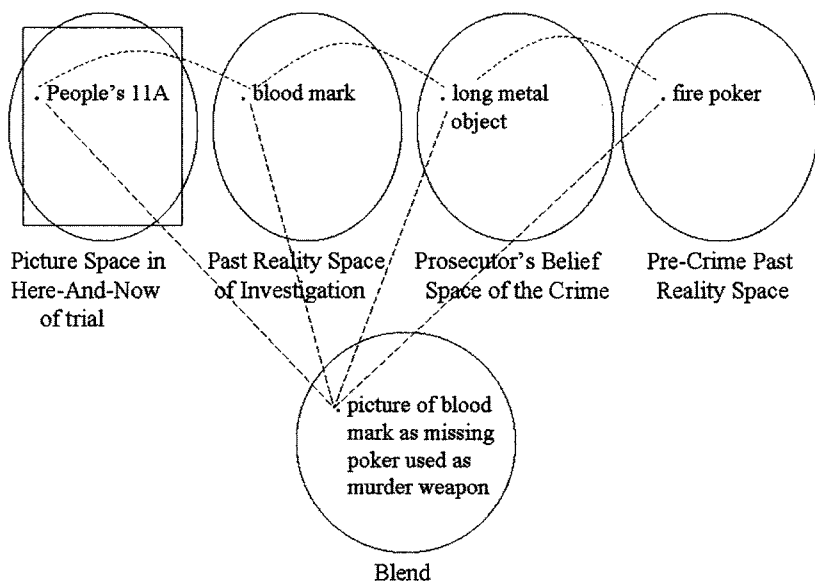


Figure 4. The picture of a blood mark as the murder weapon

Belief space of the crime, and the fireplace poker in the Past Reality space preceding the crime (see Figure 4).


Once the weapon has been set in the scene through this sequence of conceptual operations, it can be presented as operating in the prosecutor's (Professional) Belief space of the crime. See for instance the extract below (Vol. 6, 1373:27-1374: 5):

- (13) he takes a weapon like the fireplace poker, the weapon that made the red mark on the bedspread, and he takes that and hits her right across the head. Now, stop and break this down for a minute. What else could your intent be when somebody takes a fireplace poker and hits somebody across the head?

Note that the weapon is first tentatively presented as “a weapon like a fireplace poker,” which is then once more explicitly mapped onto the element that left the blood mark in the Picture space. Since in (13) it had been established that the weapon most probably corresponded to the fireplace poker in the Past Reality space prior to the crime, that is, the Past Reality space prior to the poker's mysterious disappearance, the prosecutor can subsequently speak of the weapon directly as “a fireplace poker”. Apparently, the prosecutor was so convinced of the identity connection between the missing fireplace tools and the murder weapon that he had asked to bring a similar type of poker to the courtroom for demonstrative purposes. The defense had objected to this, arguing that it was not “100% sure” that the “heavy blunt instrument” used was a fireplace poker, as “there was no instrument that was admitted into evidence” (Int.9-DC: 34). Interestingly, in order to illustrate the manner of killing all the same, the prosecutor used a wooden pointer as an alternative.

Consider the fragment below (Vol. 6, 1455:10–22) and the corresponding images borrowed from KSNB television of the prosecutor delivering this piece of discourse:

- (14) As he batters her with the weapon, [slow battering movements with pointer, while looking at jury, pic. 1] if he's even hitting just up here [touches back of pointer] or somewhere along the length of the weapon, [smooth movements along pointer, pic. 2] centrifugal force is going to keep directing the blood all the way to the end [touches end of pointer]. That's why he [blood spatter expert] said you'd expect to see some blood even on the back of the shirt [turns around] of the person who is swinging the fireplace tool [points to own back, pic. 3], because the blood is going to come off the end [touches end of pointer] [...] So as he casts it down after it's wet, blood flies to the end [touches end of pointer]



picture 1

picture 2

picture 3

Here, the wooden pointer stands for the missing weapon. Thus, the pointer serves as a material anchor to help interpret the prosecutor's regular movements in the Here-and-Now space as corresponding to movements in the Past Reality space of the crime, conceptualized within the (Professional) Belief space of the prosecution team. The prosecutor's intentional hand movements along the pointer stand for the blood's path due to gravity. At the same time, the prosecutor uses his own body in the Real space of the courtroom as a surrogate, as it stands for the attacker's body in the crime scene. Clearly, the prosecutor's mimicking gestures are not merely an embellishment of speech, but integral parts of his discourse production, which is consistent with recent studies on gestures accompanying language (McNeill 2000, 2005; Enfield 2003). Note too that the prosecutor's movements are slow and that he accompanies them with an explanation and eye-contact with the jury. This resembles more a teacher's demonstration than a realistic reconstruction of a fatal attack. This indicates yet another content-context blend, in which the pragmatic and the metapragmatic level become integrated with each other. The prosecutor is dramatizing his description, that is, illustrating the events dealt with, as well as providing a commentary on these events for the sake of the audience in the Here-and-Now space.⁴

4. Not surprisingly, the integration of description and demonstration, involving the partitioning off of body parts, allowing the presentation of the scene narrated from different viewpoints, also occurs in sign language (cf. Dudis 2004a; Liddell 2003).

5. Empathy and the generic 'you'

Besides the content-context blends that the writer and the prosecutor explicitly set up in their discourses, I suggest that they both also succeeded in implicitly prompting an identity blend between the addressees (readers and jury members) and the discourse characters (victim and defendant). I believe that this conceptual operation reflects the universal cognitive capacity to put oneself in someone else's shoes, which is fundamental to the experience of sympathy. In my feedback interviews with trial participants and attendees in the case at hand, for instance, interviewees often expressed their feelings about the case through drawing analogies with their own lives, as well as with how they believed they would feel – or *I*, their interlocutor, would feel – under the same circumstances as the individuals talked about. For example, the writer who wrote the manuscript and the juror interviewed – both females – tried to explain for themselves why the victim neither divorced the defendant nor suspected any malicious intent on his part by blending themselves and their (ex-)husbands with the couple. The writer even (rhetorically) put both herself and me – her interlocutor – in the shoes of an average victim and assaulter, thereby blending the roles of the discourse characters in the Verbal space with the interlocutors in the Here-and-Now space of the interview (Int.8-Nov: 19):

- (15) They also found a very high correlation between injuries in the face and a very *personal* relationship of hatred towards the victim. For instance, if you and I want to just kill each other, or just kill somebody, we wouldn't necessarily hit them 20 times in the head!

Here, the writer explains the results of a study showing a correlation between the number of injuries to the face of a victim and the attacker's amount of hatred towards that victim. This was relevant to the case at hand, in which the victim's face had been severely battered. The illustration of this point with an imaginary crime committed in which we, the two communicative participants, are both victims and attackers should be regarded as quite extraordinary if one bears in mind that we were strangers engaged in an interview which was being tape-recorded for subsequent analysis. Still, it seems that structuring the content of the Verbal space in terms of the Here-and-Now space was more important than truth, relevance or conversational etiquette. More generally, I believe that this constitutes a vivid exploitation of a common – and perhaps even universal – blending type. Critically, I propose that this blending type does not only become manifest in discourse, or in the grammar of American Sign Language (Liddell 1995, 1998), but also in the semantics of the generic 'you.' Consider in this regard the extract below, also from my interview with the writer (Int.8-Nov: 14, 15):

- (16) If you are a witness and you lie about something people assume that you are lying about everything else [...] That's the way the trial system works. [...] the idea being that if you and I are in the jury and we see him lie about one thing, then we assume they are lying about everything else [...] they try to trap you into a statement, and then impeach you with your own testimony [...] So you'll say, 'no I wasn't at the store at eight o'clock,' and I take out the testimony and I said 'aha! you testify here that you were at the store at eight o'clock!'

In this extract, the writer explains the working of the trial procedure, using the second person pronoun 'you' to refer to both a witness and a jury member, while the pronoun 'I' serves to refer first to a jury member and then to a lawyer. The two pronouns are used in order to illustrate a general scenario, and thus do not refer to the addresser (the writer) or the addressee (myself) in the research interview. Rather, they are to be interpreted as referring to a generic witness, jury member, and attorney. Thus, roles in the Verbal space are dynamically filled with values in the Here-and-Now space, even though they still need to be understood as roles. Note that in the cases in which the second person pronoun is used, it would be hard to find a clear-cut distinction between this use and that of the generic 'you.' Indeed, it would be difficult to argue that the generic 'you,' used repeatedly by my informants, sets up a basic conceptual configuration that differs significantly from an identity blend between individuals in the Here-and-Now and individuals in the Verbal space. It seems more accurate to state that there exists a continuum between the explicit blending of addressees with discourse characters, as in the Bypass blend or in (15) and (16), and the implicit blending through the use of a generic 'you.'

6. Summary and conclusions

This paper presented an analysis of a writer's description of a crime compared with a prosecutor's description of that same crime in his closing argument to the jury.

I hope to have shown that the pieces of discourse dealt with involved a conceptual link between (at least) the supposed Past Reality space of the crime and the ongoing communication in the Here-and-Now space. This link facilitated the creation of narratives that were halfway between fact and fiction. Sequential stories were constructed out of the presentation of fragmented events occurring in different temporal and spatial realms, some of which might have no counterparts in actuality. Spectacular and unusual as this rhetorical device may seem, its basic underlying conceptual operations are extremely common in litigation and criminal narrative. (At least) in judicial argumentation, this technique is often used because it allows one to simultaneously attend to all the relevant concerns in the same conceptual domain (cf. Pascual 2002; Coulson and Pascual 2006). In fact, it does not even seem possible to understand, reason, or talk about past events with which one has had no direct experience without evoking underlying conceptual operations of this kind. In particular, the data analyzed seems to show that the overall conceptual configuration underlying a piece of narrative such as a non-fiction book may be significantly less complex than its counterpart in an actual prosecutor's discourse in a high-profile criminal trial. This indicates that fictivity and imagination do not only play an important role in literature, but also in legal language, possibly precisely when the stakes are particularly high (cf. Coulson and Pascual 2006). This is non-trivial, since it is generally accepted that evidential law is "conceptualized as organized around facilitating the presentation and contestation of what happened, of 'facts' and 'the truth'" (Philips 1992: 250).

At the same time, in the two pieces of discourse discussed, the narrator conceptually turned addressees into direct 'observers' of the story, in such a way that communicative context and discourse content became one. Even though this occurred only implicitly

and for the sake of argument, it successfully managed to achieve sympathy. Indeed, the conceptual integration of communicative context and discourse content is critical for gaining sympathy and cognitive involvement from the audience. It may be postulated that one cannot empathize with someone else without mentally engaging in an identity blend with that individual. In fact, it seems reasonable to me to postulate that an identity blend between the interlocutors in the situation of communication (with their own past experiences and expectations of future ones) and the characters and events in the Verbal space may also be critical to ordinary comprehension and reasoning. This idea is consistent with work on simulation semantics (cf. Glenberg and Kaschak 2002; Matlock 2004; Bergen and Chang 2005; Feldman 2006), which shows that language understanding critically involves mental simulation of linguistic content. In other words, it is postulated that in order to produce or understand meaningful language, language users mentally imagine themselves perceiving or enacting the content of an utterance or piece of discourse. Along the same vein, I have suggested that the generic 'you' sets up a basic conceptual configuration based on an identity blend between the interlocutor(s) in the Here-and-Now space and the role(s) and/or the character(s) in the Verbal space.

In short, in the examples discussed in this chapter, the recognition of the context of communication is essential to constructing the intended meaning. I believe this constitutes the norm rather than the exception in ordinary language use. Bearing this in mind, I propose that the Current Discourse Space constitutes a blended space of the Verbal and the Here-and-Now spaces. To put it differently, I suggest that the conceptual domain shared by those engaged in communication needs to be understood as resulting from the blending of the configuration set up by the linguistic input with the input coming from the situated context of communication. This means that the Base space, i.e. the "starting point for the construction to which it is always possible to return" (Fauconnier 1994: xxii), always corresponds to the Here-and-Now space. This idea is consistent with the conceptual blending model developed by the Aarhus school, in which the situation of cognizing constitutes the ground upon which space building occurs (Brandt and Brandt 2005).

This paper is thus a call for an approach to blending and cognitive linguistics in general which takes into account both the content and the context of language production and interpretation. Indeed, my work is in tune with a view of language as essentially context-bound and interactively organized (cf. Cicourel 1973; Duranti and Goodwin 1992). My work is hence also in line with cognitive linguistic approaches that regard intersubjectivity as a fundamental dimension of linguistic meaning as well as grammar (cf. Sinha 1999; Verhagen 2005; Janssen 2007). Lastly, I believe the paper also shows that understanding discourse and communicative phenomena can help us better understand linguistic phenomena. It seems therefore useful to study linguistic constructions within a broad, situated discourse context.

Index

- A**
analogy 17, 22, 23, 24, 147, 148, 158, 159
anaphora 9, 38, 441, 442, 443, 444, 445, 446, 447, 448, 455, 458, 463
- B**
backgrounding 8, 331, 332, 356, 372
behavioral profile 60, 63, 64, 65, 68, 69, 70, 72, 73, 74, 100
biocultural niche 7, 289, 294, 300, 301, 302, 306, 307
Blending Theory 10
blend 10, 151, 152, 154, 156, 157, 158, 163, 164, 165, 167, 170, 172, 173, 174, 178, 179, 226, 230, 231, 465, 467, 468, 475, 485, 487, 491, 493, 501, 503, 504, 506, 507, 508, 511, 512, 513, 514, 515, 516
compression 5, 148, 149, 151, 156, 477, 484, 485, 489, 498, 508
conceptual blending 147, 158, 198, 465, 500, 501, 515
conceptual integration 4, 5, 6, 149, 150, 159, 161, 162, 225, 226, 468, 476, 499, 501, 503, 507, 514
Conceptual Integration Theory 500
constructional blend 5, 161, 169, 180
double-scope blend 5, 148
grounded blend 501, 506
integration network 147, 148, 150, 156, 158, 159, 179, 266, 475, 515
mirror blend 5, 148
simplex blend 5, 148
- C**
cluster analysis 65, 67, 69, 73
coevolution 293, 295, 296, 297, 298, 307
cognitive grammar 63, 101, 207, 267, 279, 280, 304, 463, 516
cognitive operations 147, 148, 149, 336
cognitive processes 8, 79, 107, 297, 312, 323, 371, 397, 421
cognitive science 1, 8, 26, 143, 312, 371, 422, 437
cognitive semantics 4, 74, 102, 396, 397, 412
communicative context 501, 505, 514
composition 6, 22, 69, 162, 181, 191, 226, 231, 233, 235, 240, 243, 246, 258, 265, 423, 436
compositionality 19, 23, 24, 179, 180, 187, 238, 270, 281
conceptual category 8, 331
Conceptual Metaphor Theory 5, 127, 128, 129, 130, 135, 136, 137, 138, 140, 141, 142
conceptual metaphor 4, 5, 101, 112, 127, 130, 131, 140, 143, 168, 240, 246, 326, 465, 498
metaphor 2, 4, 6, 11, 96, 105, 107, 108, 109, 110, 111, 112, 114, 115, 116, 117, 123, 125, 127, 128, 129, 130, 135, 136, 140, 144, 147, 148, 149, 150, 151, 157, 158, 159, 167, 180, 205, 225, 226, 227, 229, 231, 239, 240, 246, 248, 265, 324, 326, 327, 332, 408, 465, 498, 503, 505, 515
metaphorical interpretation 4, 108, 109, 113, 114, 116, 117, 121
metaphorical mapping 4, 109, 148
metaphorical meaning 107, 109, 110, 111, 114, 115, 116, 117, 118, 121, 123, 124
metaphorical source 128
spatial metaphor 130, 140, 142, 143
conceptual projection 10, 128, 145, 465
conceptual reference point 9, 447
construction grammar 101, 185, 192, 197, 198, 266, 267, 396
constructional compositionality 6, 161, 165, 173
constructional meaning 6, 7, 102, 161, 163, 177, 191, 225, 235, 248, 253, 255, 258, 262, 267
co-occurrence 48, 59, 60, 61, 62, 63, 64, 65, 83, 95, 97
corpus linguistics 59, 82, 101, 102, 103, 185
correspondence analysis 4
culture 7, 27, 81, 101, 150, 271, 289, 290, 291, 292, 293, 295, 296, 297, 298, 300, 302, 306, 307, 308, 402, 448, 466, 470
- D**
Darwinism 295, 300
definiteness 18, 26, 276, 462
derivation 6, 234, 235, 243, 246, 247, 260, 264, 265, 277, 355
directionality 8, 226, 229, 297, 331, 366, 372
discourse context 4, 10, 38, 107, 113, 117, 124, 442, 456, 485, 514
- E**
ecological niche 291, 293, 294
embodiment 2, 7, 8, 287, 311, 312, 313, 314, 327, 328, 329, 368

- Lakoff, G. & Johnson, M. 1999. *Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought*. New York: Basic Books.
- Lakoff, G. & Turner, M. 1989. *More than Cool Reason: A Field Guide to Poetic Metaphor*. Chicago: University of Chicago Press.
- Langacker, R. W. 1999. Virtual reality. *Studies in Linguistic Sciences* 29 (2): 77–103.
- Langacker, R. W. 2001. Discourse in cognitive grammar. *Cognitive Linguistics* 12 (2): 143–188.
- Liddell, S. K. 1995. Real, surrogate and token space: Grammatical consequences in ASL. In *Sign, Gesture and Space*, K. Emmorey & J. Reilly (eds), 19–41. New Jersey: Lawrence Erlbaum.
- Liddell, S. K. 1998. Grounded blends, gestures and conceptual shifts. *Cognitive Linguistics* 9 (3): 283–314.
- Liddell, S. K. 2003. *Grammar, Gesture, and Meaning in American Sign Language*. Cambridge: Cambridge University Press.
- Luchjenbroers, J. 2006. Discourse, gesture, and mental spaces manoeuvres: Inside versus outside F-space. In *Cognitive Linguistics Investigations: Across Languages, Fields and Philosophical Boundaries*, J. Luchjenbroers (ed.), 87–105. Amsterdam: John Benjamins.
- Matlock, T. 2004. Fictive motion as cognitive simulation. *Memory and Cognition* 32: 1389–1400.
- McNeill, D. (ed.). 2000. *Language and Gesture*. Cambridge: Cambridge University Press.
- McNeill, D. 2005. *Gesture and Thought*. Chicago: University of Chicago Press.
- Oakley, T. 1998. Conceptual blending, narrative discourse, and rhetoric. *Cognitive Linguistics* 9 (4): 320–360.
- Panther, K.-U. & Thornburg, L. L. 2000. The EFFECT FOR CAUSE Metonymy in English Grammar. In *Metaphor and Metonymy at the Crossroads*, A. Barcelona (ed.), 215–232. Berlin: Mouton de Gruyter.
- Pascual, E. 2002. *Imaginary Dialogues: Conceptual Blending and Fictive Interaction in Criminal Courts*. Utrecht: LOT 68.
- Pascual, E. 2008. Text for context, trial for trialogue: An ethnographic study of a fictive interaction blend. *Annual Review of Cognitive Linguistics* 6: 50–82.
- Philips, S. 1992. Evidentiary standards for American trials: Just the facts. In *Responsibility and Evidence in Oral Discourse*, J. H. Hill & J. T. Irvine (eds), 248–259. Cambridge: Cambridge University Press.
- Sinha, C. 1999. Grounding, mapping, and acts of meaning. In *Cognitive Linguistics: Foundations, Scope, and Methodology*, Th. Janssen & G. Redeker (eds), 223–255. Berlin: Mouton de Gruyter.
- Sinha, C. 2005. Blending out of the background: Play, props and staging in the material world. *Journal of Pragmatics* 37 (10): 1537–1554.
- Tobin, V. 2006. Joint attention in reading communities: Overhearer effects, routinization, and reading Sherlock Holmes. Paper presented at the *Conceptual Structure, Discourse and Language Conference*. University of California, San Diego, November 3–5, 2006.
- Verhagen, A. 2005. *Constructions of Intersubjectivity: Discourse, Syntax, and Cognition*. Oxford: Oxford University Press.
- Williams, R. F. 2004. *Making Meaning from a Clock: Material Artifacts and Conceptual Blending in Time-Telling Instruction*. PhD Dissertation. University of California, San Diego.

- P**
- partonomy 7, 269, 270, 272, 280, 282, 283
- perception 3, 8, 27, 33, 42, 62, 71, 229, 280, 281, 290, 294, 299, 312, 313, 328, 332, 397, 437, 457, 482, 484, 485, 501, 504, 505
- phenotype 293, 294, 295, 298
- phylogenesis 299
- pragmatics 9, 61, 304, 305, 395, 441
- prescriptive grammar 15
- principled polysemy 4, 58, 60, 63, 78, 79, 80, 99
- procedural 17, 18, 23, 24
- process and product 17
- prompts 15, 18, 163, 164, 168, 172, 173, 175, 256, 468
- proprioception 8, 30, 31, 33, 41, 42, 313, 321, 323, 332
- prototype 4, 8, 70, 80, 100, 230, 231, 239, 331, 346, 348, 360
- prototypicality 57, 64, 73, 74, 101, 108
- psychology 1, 145, 294, 327, 332, 396, 397, 421
- R**
- radial categories 3, 73, 103
- Real space 501, 505, 506, 508, 509, 511
- reference point 5, 9, 161, 174, 176, 178, 258, 346, 441, 442, 443, 444, 445, 446, 447, 448, 449, 450, 451, 452, 454, 456, 457, 458, 459, 462, 463
- Relevance Theory 17
- S**
- semantic pole 2, 16, 27, 233, 240, 243, 258, 262
- semiotics 144
- sense distinction 78, 79
- sexual selection 292, 294, 295, 299
- sign 16, 20, 52, 151, 152, 280, 302, 303, 309, 403, 411, 412, 469, 473, 482, 511
- signification 302, 303
- similarity 5, 57, 58, 59, 65, 66, 69, 95, 109, 128, 133, 134, 135, 136, 137, 138, 139, 140, 141, 142, 143, 144, 145, 172, 210, 217, 219, 250, 374, 376, 377, 433
- social cognitive linguistics 395, 412, 413, 417, 418
- social fact 300, 301, 302, 305, 306, 404
- social institution 7, 8, 289, 300, 306, 307
- spatial concept 331, 347
- structuralism 15
- symbol 42, 52, 53, 208, 247, 299, 313, 316, 327, 328, 337, 403, 411, 490, 493
- sympathy 324, 503, 512, 514
- synonymy 4, 53, 57, 60, 69, 101, 104, 414
- syntax 4, 7, 26, 60, 72, 73, 77, 81, 164, 199, 218, 225, 245, 262, 265, 269, 270, 271, 273, 275, 279, 282, 284, 285, 312, 323, 346, 361, 367, 395, 396, 423, 424, 425, 426, 427, 434, 435, 436, 437, 440, 441
- T**
- tacit knowledge 19
- temporal viewpoint 8, 331
- U**
- usage event 3, 16, 20, 22, 23, 57, 168, 181
- usage fundamentalism 2, 16
- usage-based 4, 6, 15, 16, 19, 22, 23, 25, 59, 60, 72, 74, 77, 80, 81, 83, 84, 99, 101, 102, 185, 186, 198, 199, 266, 267, 329, 435
- W**
- word meaning 19, 22, 23, 75, 127, 135, 180, 204, 219, 366, 396

- emergence 3, 7, 22, 27, 51, 154,
162, 179, 263, 297, 299, 306, 316,
332, 354, 365, 438
- encyclopaedic knowledge 99
- encyclopaedic semantics 77,
81, 88
- environmentalism 291
- epigenesis 297, 298, 299, 300,
306, 308
- experimental methods 8, 58
- expository discourse 422, 423,
429, 436, 437, 439
- F**
- foregrounding 8, 115, 331, 332,
364, 372
- formal semantics 331, 396, 397
- framing 10, 47, 147, 148, 149, 169,
171, 175, 177, 178, 179, 372,
373, 376, 389, 397, 443, 447,
450, 462, 477, 478, 479, 505
- frame metonymy 6, 161, 162,
164, 169, 178
- G**
- generative grammar 15, 277, 278,
279, 396, 397
- genitive 5, 6, 11, 161, 163, 165, 170,
171, 172, 173, 174, 175, 176, 177,
178, 179, 181
- genotype 292, 295, 296, 298
- global text quality 9, 424, 427,
428, 432, 433, 434, 435, 437
- Graded Salience Hypothesis 4,
11, 107, 108, 109, 110, 111, 112, 113,
117, 123, 125
- H**
- Here-and-Now space 501, 502,
503, 506, 511, 512, 513, 514
- I**
- ID tag 59, 60, 61, 62, 63, 64, 65,
66, 67, 68, 69, 73
- idioms 189, 193, 195, 197, 198,
322, 326
- image schema 8, 57, 312, 314, 321,
322, 323, 325, 326, 327, 329, 331,
332, 333, 334, 335, 336, 337, 339,
342, 346, 348, 349, 354, 356, 360,
366, 368
- indeterminacy 21, 22, 187, 415,
416, 418
- indeterminate 20, 21, 22, 418
- inference 8, 129, 141, 142, 145,
305, 345, 349, 351, 361, 383, 387,
388, 483
- inheritance 6, 185, 187, 196, 198,
209, 226, 247, 248, 265, 473
- instructional semantics 15
- inter-subjectivity 10
- L**
- language development 6, 9, 203,
312, 440
- language use 2, 7, 9, 11, 15, 17,
18, 19, 22, 23, 24, 29, 52, 74, 77,
80, 103, 113, 132, 202, 203, 267,
296, 304, 307, 327, 416, 417, 418,
437, 514
- LCCM Theory 3, 11, 27, 28, 31, 32,
36, 40, 51, 53
- cognitive model 3, 10, 18, 25,
27, 28, 30, 31, 40, 41, 49, 50,
52, 53
- lexical concept 3, 10, 25, 27,
28, 30, 31, 32, 33, 36, 37, 38,
39, 40, 51, 52, 53
- lexical representation 31
- semantic potential 32, 49,
51, 52
- semantic representation 3, 27,
28, 31, 32, 53
- semantic structure 3, 27, 31,
32, 40,
- lexical meaning 4, 18, 19, 108,
110, 116, 118, 241
- lexical network model 4, 78, 80
- Lexical representation 79, 208,
241
- lexical semantics 3, 25, 73, 82,
84, 100, 102, 104, 181, 190, 191,
266, 267
- lexical usage 9, 423, 424, 425,
435, 436
- lexicography 3, 99, 218
- linguistic metaphor 5, 127, 128,
129, 130, 131, 132, 133, 140, 142,
143
- linguistic relativity 8, 371, 372,
373, 375, 386, 389, 390, 391, 440
- M**
- mapping 99, 110, 147, 148, 151,
154, 156, 157, 158, 166, 168, 201,
226, 327, 372, 417, 420, 489, 505,
509, 516
- meaning construction 2, 27, 31,
80, 104, 107, 111, 148, 158, 360,
500
- meme 295
- memory 8, 48, 57, 135, 142, 164,
187, 207, 216, 280, 290, 374, 375,
382, 384, 385, 387, 389, 390, 396,
397, 450, 453, 463, 484, 490,
491, 503
- mental representation 1, 2, 45,
46, 58, 73, 99, 127, 128, 131, 132,
133, 141, 143, 340, 397
- Mental Spaces Theory
- mental space 17, 147, 159, 225,
226, 231, 441, 443, 445, 447,
462, 470, 477, 478, 479, 482,
483, 500, 501, 503, 516
- metonymy 101, 147, 148, 158, 161,
162, 166, 168, 172, 175, 180, 181,
225, 226, 239, 265, 465, 509
- mimesis 8, 313, 329
- modal distance 8, 331, 338
- modular 289
- modularity 292
- multifactorial method 4, 77
- N**
- nativism 290, 291
- natural selection 292
- near synonymy 3, 57, 59, 60, 74
- Neo-Darwinism 295
- network representations 3, 73
- niche construction 292, 296
- O**
- objectivism 305
- ontogenesis 300, 312
- ontology 42, 43, 44, 290, 295,
300
- organism 293, 294, 296, 297,
300, 334
- output 2, 16, 17, 18, 20, 21, 24,
25, 226

- P**
- partonomy 7, 269, 270, 272, 280, 282, 283
- perception 3, 8, 27, 33, 42, 62, 71, 229, 280, 281, 290, 294, 299, 312, 313, 328, 332, 397, 437, 457, 482, 484, 485, 501, 504, 505
- phenotype 293, 294, 295, 298
- phylogenesis 299
- pragmatics 9, 61, 304, 305, 395, 441
- prescriptive grammar 15
- principled polysemy 4, 58, 60, 63, 78, 79, 80, 99
- procedural 17, 18, 23, 24
- process and product 17
- prompts 15, 18, 163, 164, 168, 172, 173, 175, 256, 468
- proprioception 8, 30, 31, 33, 41, 42, 313, 321, 323, 332
- prototype 4, 8, 70, 80, 100, 230, 231, 239, 331, 346, 348, 360
- prototypicality 57, 64, 73, 74, 101, 108
- psychology 1, 145, 294, 327, 332, 396, 397, 421
- R**
- radial categories 3, 73, 103
- Real space 501, 505, 506, 508, 509, 511
- reference point 5, 9, 161, 174, 176, 178, 258, 346, 441, 442, 443, 444, 445, 446, 447, 448, 449, 450, 451, 452, 454, 456, 457, 458, 459, 462, 463
- Relevance Theory 17
- S**
- semantic pole 2, 16, 27, 233, 240, 243, 258, 262
- semiotics 144
- sense distinction 78, 79
- sexual selection 292, 294, 295, 299
- sign 16, 20, 52, 151, 152, 280, 302, 303, 309, 403, 411, 412, 469, 473, 482, 511
- signification 302, 303
- similarity 5, 57, 58, 59, 65, 66, 69, 95, 109, 128, 133, 134, 135, 136, 137, 138, 139, 140, 141, 142, 143, 144, 145, 172, 210, 217, 219, 250, 374, 376, 377, 433
- social cognitive linguistics 395, 412, 413, 417, 418
- social fact 300, 301, 302, 305, 306, 404
- social institution 7, 8, 289, 300, 306, 307
- spatial concept 331, 347
- structuralism 15
- symbol 42, 52, 53, 208, 247, 299, 313, 316, 327, 328, 337, 403, 411, 490, 493
- sympathy 324, 503, 512, 514
- synonymy 4, 53, 57, 60, 69, 101, 104, 414
- syntax 4, 7, 26, 60, 72, 73, 77, 81, 164, 199, 218, 225, 245, 262, 265, 269, 270, 271, 273, 275, 279, 282, 284, 285, 312, 323, 346, 361, 367, 395, 396, 423, 424, 425, 426, 427, 434, 435, 436, 437, 440, 441
- T**
- tacit knowledge 19
- temporal viewpoint 8, 331
- U**
- usage event 3, 16, 20, 22, 23, 57, 168, 181
- usage fundamentalism 2, 16
- usage-based 4, 6, 15, 16, 19, 22, 23, 25, 59, 60, 72, 74, 77, 80, 81, 83, 84, 99, 101, 102, 185, 186, 198, 199, 266, 267, 329, 435
- W**
- word meaning 19, 22, 23, 75, 127, 135, 180, 204, 219, 366, 396

In the series *Human Cognitive Processing* the following titles have been published thus far or are scheduled for publication:

- 25 **PANTHER, Klaus-Uwe, Linda L. THORNBURG and Antonio BARCELONA (eds.):** *Metonymy and Metaphor in Grammar.* xiii, 401 pp + index. *Expected July 2009*
- 24 **EVANS, Vyvyan and Stéphanie POURCEL (eds.):** *New Directions in Cognitive Linguistics.* 2009. xi, 519 pp.
- 23 **HOLŠANOVÁ, Jana:** *Discourse, Vision, and Cognition.* 2008. xiii, 202 pp.
- 22 **BERENDT, Erich A. (ed.):** *Metaphors for Learning. Cross-cultural Perspectives.* 2008. ix, 249 pp.
- 21 **AMBERBER, Mengistu (ed.):** *The Language of Memory in a Crosslinguistic Perspective.* 2007. xii, 284 pp.
- 20 **AURNAGUE, Michel, Maya HICKMANN and Laure VIEU (eds.):** *The Categorization of Spatial Entities in Language and Cognition.* 2007. viii, 371 pp.
- 19 **BENCZES, Réka:** *Creative Compounding in English. The Semantics of Metaphorical and Metonymical Noun-Noun Combinations.* 2006. xvi, 206 pp.
- 18 **GONZALEZ-MARQUEZ, Monica, Irene MITTELBERG, Seana COULSON and Michael J. SPIVEY (eds.):** *Methods in Cognitive Linguistics.* 2007. xxviii, 452 pp.
- 17 **LANGLOTZ, Andreas:** *Idiomatic Creativity. A cognitive-linguistic model of idiom-representation and idiom-variation in English.* 2006. xii, 326 pp.
- 16 **TSUR, Reuven:** 'Kubla Khan' – Poetic Structure, Hypnotic Quality and Cognitive Style. A study in mental, vocal and critical performance. 2006. xii, 252 pp.
- 15 **LUCHJENBROERS, June (ed.):** *Cognitive Linguistics Investigations. Across languages, fields and philosophical boundaries.* 2006. xiii, 334 pp.
- 14 **ITKONEN, Esa:** *Analogy as Structure and Process. Approaches in linguistics, cognitive psychology and philosophy of science.* 2005. xiv, 249 pp.
- 13 **PRANDI, Michele:** *The Building Blocks of Meaning. Ideas for a philosophical grammar.* 2004. xviii, 521 pp.
- 12 **EVANS, Vyvyan:** *The Structure of Time. Language, meaning and temporal cognition.* 2004. x, 286 pp.
- 11 **SHELLEY, Cameron:** *Multiple Analogies in Science and Philosophy.* 2003. xvi, 168 pp.
- 10 **SKOUSEN, Royal, Deryle LONSDALE and Dilworth B. PARKINSON (eds.):** *Analogical Modeling. An exemplar-based approach to language.* 2002. x, 417 pp.
- 9 **GRAUMANN, Carl Friedrich and Werner KALLMEYER (eds.):** *Perspective and Perspectivation in Discourse.* 2002. vi, 401 pp.
- 8 **SANDERS, Ted, Joost SCHILPEROORD and Wilbert SPOOREN (eds.):** *Text Representation. Linguistic and psycholinguistic aspects.* 2001. viii, 364 pp.
- 7 **SCHLESINGER, Izchak M., Tamar KEREN-PORTNOY and Tamar PARUSH:** *The Structure of Arguments.* 2001. xx, 264 pp.
- 6 **FORTESCUE, Michael:** *Pattern and Process. A Whiteheadian perspective on linguistics.* 2001. viii, 312 pp.
- 5 **NUYTS, Jan:** *Epistemic Modality, Language, and Conceptualization. A cognitive-pragmatic perspective.* 2001. xx, 429 pp.
- 4 **PANTHER, Klaus-Uwe and Günter RADDEN (eds.):** *Metonymy in Language and Thought.* 1999. vii, 410 pp.
- 3 **FUCHS, Catherine and Stéphane ROBERT (eds.):** *Language Diversity and Cognitive Representations.* 1999. x, 229 pp.
- 2 **COOPER, David L.:** *Linguistic Attractors. The cognitive dynamics of language acquisition and change.* 1999. xv, 375 pp.
- 1 **YU, Ning:** *The Contemporary Theory of Metaphor. A perspective from Chinese.* 1998. x, 278 pp.