

# The Semantics of English Prepositions

*Spatial Scenes, Embodied Meaning  
and Cognition*

---

Andrea Tyler

*Linguistics Department  
Georgetown University*

and

Vyvyan Evans

*School of Cognitive and Computing Sciences  
University of Sussex*



PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE  
The Pitt Building, Trumpington Street, Cambridge CB2 1RP, United Kingdom

CAMBRIDGE UNIVERSITY PRESS  
The Edinburgh Building, Cambridge, CB2 2RU, UK  
40 West 20th Street, New York, NY 10011-4211, USA  
477 Williamstown Road, Port Melbourne, VIC 3207, Australia  
Ruiz de Alarcón 13, 28014 Madrid, Spain  
Dock House, The Waterfront, Cape Town 8001, South Africa  
<http://www.cambridge.org>

© Andrea Tyler and Vyvyan Evans 2003

This book is in copyright. Subject to statutory exception  
and to the provisions of relevant collective licensing agreements,  
no reproduction of any part may take place without  
the written permission of Cambridge University Press.

First published 2003

Printed in the United Kingdom at the University Press, Cambridge

*Typeface* Times 10/12 pt    *System* L<sup>A</sup>T<sub>E</sub>X 2<sub>ε</sub> [TB]

*A catalogue record for this book is available from the British Library*

ISBN 0 521 81430 8 hardback

# Contents

---

<i>Preface</i>	page ix
<i>Acknowledgements</i>	xi
1 The nature of meaning	1
2 Embodied meaning and spatial experience	23
3 Towards a model of principled polysemy: spatial scenes and conceptualization	37
4 The semantic network for <i>over</i>	64
5 The vertical axis	107
6 Spatial particles of orientation	132
7 Bounded landmarks	178
8 Conclusion	229
<i>References</i>	238
<i>Index</i>	246

# 1 The nature of meaning

---

My grandmother was a great one for mixing historical lessons in with child rearing. A favorite, regularly used when one of the grandchildren was being rebuked for failing to satisfactorily complete some minor task and was, consequently, being required to do it  $over_A$ , involved pointing to the needle-point text hanging  $over_B$  the sofa which read, 'We won't come back 'til it's  $over_C$ ,  $over_D$  there.' This was inevitably followed by the question, 'Where would the world be if they hadn't done their jobs properly?'

This text might strike contemporary readers as a little unusual. References to lines from old war songs and needle-pointed mottoes hanging on the wall belong to a bygone era and anecdotes relying on them are likely to be somewhat vague. However, what is even more striking about the text is exactly what typical native speakers of English are likely to find unremarkable, namely the numerous, very different interpretations assigned to the single word, *over*. In this short text, *over* has four distinct interpretations –  $over_A$  can be paraphrased by 'again',  $over_B$  by 'above',  $over_C$  by 'finished' and  $over_D$  by 'in some other place'. For us, the fundamental question that texts such as the one above raise is whether the various meanings regularly associated with a single word are simply accidental (the fact that *over* has four very different meanings might, after all, be a bizarre accident), or systematically related.

Linguists have often assumed that words constitute lexical forms that are conventionally paired with meanings, and that these form-meaning pairings are stored in a mental dictionary or lexicon. Traditional approaches to the mental lexicon have tended either to ignore the issue of whether distinct meanings associated with a single form are related or to assume that the relationships are arbitrary, which is to say, the forms are unrelated. However, the linguist Bernd Heine (1997) has observed that finding a satisfactory solution to the problem of how to represent the multiple meanings associated with a single linguistic form is both a central and a controversial issue for linguistic theory. The position taken on this question affects not only how we model the semantics of individual lexical items and the architecture of the mental lexicon, but also the rest of one's model of language.

The lexicon represents a pivotal interface between syntax, semantics and pragmatics; the representation of the semantic component of lexical items has crucial implications not only for a theory of word meaning but also for a theory of sentence-level meaning construction. At stake are issues concerning the source of the information that is necessary in the interpretation of an utterance and the appropriate location of the productive (rule-governed) elements of the linguistic system. Such issues bear on the interaction between words and the human conceptual system. In addition, establishing the semantic content of the lexical representations directly impinges on the distinction between conventionalized linguistic knowledge and encyclopedic, general world knowledge in the process of meaning construction, which is to say, the traditional distinction between pragmatics and semantics.

In this book we take up the challenge of how best to represent the distinct meanings or senses associated with a single lexical form. We do so through an examination of the semantics of a range of English spatial particles, such as *over*, *up*, *down*, *in* and *out*, etc. There are a number of reasons for choosing spatial particles. Perhaps most importantly the variety and complexity of the numerous different meanings associated with even a single spatial particle represents a significant descriptive challenge. Hence, insights gleaned from such an analysis promise to have considerable applicability to other classes of words.<sup>1</sup> But also important is the relatively transparent experiential basis of spatial morphemes. The meanings of this set of words are clearly grounded at some level in our spatio-physical interaction with the world. Hence, investigating the meanings associated with spatial particles will offer fundamental insights into the relation between language, mental representation and human experience.

Our investigation leads us to the conclusion that the various meanings associated with spatial particles are related in systematic and highly motivated ways. In other words, we advance a polysemy approach to word meaning (our polysemy commitment), arguing that the multiple, distinct meanings associated with the same lexical form are often related.<sup>2</sup> We suggest that the distinct but related senses associated with a single spatial particle constitute a semantic

<sup>1</sup> Evans (2000) has successfully extended the same model to an examination of the conceptualization of time.

<sup>2</sup> It is important to emphasize that we do not claim that all representations of a form in semantic memory are part of that form's semantic network. We adhere to the basic tenet of cognitive linguistics that language is usage based and, as a result, likely to be more redundant than traditional accounts. As such, collocations and longer phrases involving context may become established in semantic memory. Thus, phrases such as *over the bridge* or perhaps constructions (Goldberg, 1995) such as 'Motion Verb + *over the bridge*', which appear largely decomposable, may be part of permanent semantic memory. The construction itself, then, may take on additional meaning. For the most part, it is beyond the scope of this book to examine the many longer phrases and possible constructions in which English spatial particles regularly participate.

network organized with respect to a primary sense. In chapter 3, where we advance the model upon which the analysis in this book will be based, we explore in detail what constitutes a primary sense, and in what way other senses might be diachronically and perhaps developmentally related to this sense. Each distinct sense is potentially subject to a number of inferencing strategies which account for additional or on-line interpretations. Consequently, in the course of developing a theory of word meaning and mental representation, we advance a concomitant theory of meaning construction or conceptual integration. Our findings reveal the largely non-idiosyncratic and systematic organization of the mental lexicon and the highly creative nature of the human conceptual system.

A number of basic assumptions underlie our approach:

- Language (lexical items and the syntactic arrangements in which they occur) radically underdetermines the rich interpretations regularly assigned to naturally occurring utterances. A consequence of this is the assumption that lexical entries, albeit crucial, act merely as prompts for meaning construction, and that meaning construction is largely a conceptual process, involving elaboration and integration of linguistic and non-linguistic information in a highly creative way (Fauconnier, 1994, 1997; Fauconnier and Turner, 1998, 2002; Turner, 1991, 1996). This is discussed later in this chapter and in chapter 3.<sup>3</sup>
- The representation of meaning is fundamentally conceptual in nature. Language does not refer directly to the ‘real world’. Rather, language refers to what is represented in the human conceptual system. The conceptual system contains conceptual structure (i.e., concepts, schemas, scripts, etc.) which indirectly reflects and interprets the world as mediated by human experience and perception (Fauconnier, 1997; Jackendoff, 1983, 1987, 1990, 1992; Langacker, 1987, 1991b). This is discussed in detail later in this chapter.
- Conceptual structure is a product of how we as human beings experience and interact with the spatio-physical world we inhabit. The world ‘out there’ provides much of the raw sense-perceptual substrate for the conceptual system. However, how and what we experience is crucially mediated by the precise nature of our bodies and our unique neuro-anatomical architecture. In other words, experience is embodied (Johnson, 1987; Lakoff, 1987; Lakoff and Johnson, 1999; Mandler, 1992, 1996; Sweetser, 1990; Varela, Thompson and Rosch, 1991). This is the subject of chapter 2.
- Language is a continually evolving, organic system. Hence, to study the synchronic ‘slice’ of a language will reveal only one point in a continuum of change (Bybee *et al.*, 1994; Hopper and Traugott, 1993). Synchronic studies,

<sup>3</sup> We specifically attempt to build this assumption into our model through the inferencing strategy of real-world force dynamics.

such as the present one, must be mindful that lexical structure of even a single form (its semantic network) will exhibit the co-existing ‘layers’ of its past.<sup>4</sup>

- The development and extension of lexical meaning result from pragmatic inferencing (i.e., situated implicatures). This leads to conceptual reanalysis and concomitant conventionalization of the inference as a new meaning component associated with the linguistic form. This results in the development of a semantic network. Borrowing terminology from Traugott (e.g., 1989), we refer to this context-based process of lexical meaning extension as *pragmatic strengthening*. Hence, meaning extension is usage based and pragmatic in nature. This is discussed in chapter 3.<sup>5</sup>

### **Approaches to the representation of distinct meanings associated with a single form**

The question of how best to model the distinct meanings associated with a single lexical form has been approached from three perspectives. These are homonymy, monosemy and polysemy. We turn now to a consideration of these three approaches.

In the text that began this chapter, the form *over* is associated with four different meanings. We paraphrased these four meanings as ‘again’, ‘above’, ‘finished’ and ‘in some other place’ respectively. In attempting to account for these different meanings, one could assume that they are unrelated. That is, one could argue that as speakers of English we have simply memorized several

<sup>4</sup> We attempt to build this assumption into our model through our choice of the primary sense associated with each spatial particle, which reflects both the diachronic and ontogenetic nature of the semantic network. The choice of the appropriate primary sense for a spatial particle has been a controversial one in discussions of the semantic networks of spatial particles (e.g., Dewell, 1994; Kreitzer, 1997; Lakoff, 1987). Some scholars in the field have openly appealed to intuition concerning the most primary meaning of a spatial particle (Dewell, 1994). However, intuitions on ‘the’ primary sense of many spatial particles vary widely. Others (e.g., Lakoff) have argued for a primary sense that best fits a particular analysis of a particular preposition. Our goal is to begin to work out a principled framework that accounts for the development of the polysemy networks for all English spatial particles. Working out principles and criteria which apply to the entire system of particles places substantial constraints on the nature of the primary sense which are not apparent when one focuses on the analysis of a single spatial particle in isolation. Our study of twenty spatial particles has led us to conclude that positing a proto-spatial scene (which includes both a configurational component and a functional component, to be outlined in detail in chapter 3) for the primary sense (the proto-scene) allows us to develop a consistent, principled analysis which calls for a minimal amount of theoretical machinery. Consultation of the OED has also revealed that the proto-scenes for each spatial particle also tend to represent the diachronically earliest uses of the lexical form. Moreover, in attempting to explain the relationships among senses associated with a single form, we found ourselves explaining how attested uses plausibly developed from prior uses. Again, we have attempted to constrain these arguments for a plausible path of development to a minimal number of theoretical constructs which would apply to all spatial particles.

<sup>5</sup> We attempt to build this assumption into our model through the notion of pragmatic strengthening (see chapter 3).

distinct meanings which are coded by the form *over*. This might be viewed as being parallel to our learning that the form *bank* is arbitrarily associated with both 'the sides of a river' and 'a certain type of financial institution'. Accordingly, it would be claimed that there are several distinct form-meaning pairings for *over* which language users represent in their mental lexicons. This position would thus posit that each of the form-meaning lexical entries are homonyms, which is to say that they are unrelated. The fact that the different senses are coded by the same linguistic form is presumably just an accident. This is essentially the position taken by traditional representations of the lexicon. Starting as early as Bloomfield (1933) and rearticulated as recently as Chomsky (1995), influential linguistic theories have asserted that the lexicon is the repository for the arbitrary and the idiosyncratic. Such analyses hold that all creativity and systematicity is in the morpho-syntactic component.

The homonymy approach suffers from a number of weaknesses when we attempt to account for words such as *over*. First, it ignores any systematic relationships among the distinct meanings associated with a single linguistic form. This stands in sharp contrast to a growing body of work (e.g., Brugman and Lakoff, 1988; Jackendoff, 1997; Lakoff, 1987; Langacker, 1987, 1991b; Levin, 1993; Lindner, 1981; Pustejovsky, 1998) which demonstrates that systematic, rule governed relationships do exist in the lexicon.

Second, the homonymy position takes a narrow synchronic view. That is, it fails to represent language as an evolving system whose changes over time are largely constrained in a motivated, principled manner. The synchronic semantic network associated with a lexical item is a historical product. In assuming that distinct meanings within a semantic network are arbitrarily related, the homonymy approach makes the implicit claim that the process of meaning extension itself is arbitrary, leading to the unsatisfactory conclusion that language change is ad hoc, lacking motivation. This contradicts the view that language evolution is a systematic process, as revealed by the voluminous grammaticalization literature (e.g., Bybee *et al.*, 1994; Heine *et al.*, 1991; Hopper and Traugott, 1993, for overviews, summaries and references).

Moreover, it is reasonable to assume that at an earlier stage in the language, a form such as *over* had fewer distinct, conventionalized meanings associated with it;<sup>6</sup> thus, many of the uses now conventionally associated with the form at one point represented novel uses. The homonymy approach begs the question of why it should be the case that a speaker would choose to use a particular established form in a novel way, rather than coining a new phonological string altogether.

It is perhaps self-evident that an important function of language is communication. Moreover, communication is fundamentally purposeful (Gumperz,

<sup>6</sup> This point is also made by Sweetser (1990).

1982). This fact places certain non-trivial constraints on the use of lexical items. It is obvious that a speaker intending to communicate, and hence achieve the desired purpose, would not use a lexical form with one established meaning to indicate something else, unless the speaker assumed the listener could readily work out the novel usage. In order for a novel use to be readily interpretable by the hearer, meaning extension must be somehow constrained and systematic. This strongly suggests that when a speaker uses a form with an established meaning to indicate something other than the conventional meaning, the choice of which lexical item to select is motivated. If this were otherwise, the speaker could not assume that the listener had a reasonable chance of interpreting the novel use. This line of reasoning suggests that there must be something about the conventional meaning associated with the lexical item that led the speaker to choose that lexical form rather than some other.

Finally, the homonymy approach fails to explain the ubiquity of the phenomenon. Every spatial particle of English demonstrates multiple senses. Moreover, careful examination of spatial particles as a word class reveals regular patterns of meanings across the individual members of the class. For instance, all spatial particles whose primary senses have a distinct goal sense share a clearly defined set of properties. The homonymy approach argues that this too is accidental. While we readily acknowledge that the ‘accidents’ of history have resulted in instances of homonymy, the failure to account for the considerable systematicity that does exist misses important generalizations. As the homonymy approach fails to recognize that distinct meanings may be motivated and, hence, at some level systematically related, we are forced to conclude that it is inadequate.

An alternative approach, monosemy, has been advocated by Charles Ruhl (1989). Ruhl argues in detail that forms are paired with a single highly abstract meaning. This abstract meaning can be filled in by contextual knowledge, such that all the distinct meanings associated with a lexeme are derived. This position is termed monosemy, as it holds that the multiple meanings associated with a particular form are merely contextually derived variants of a single monosemous meaning.

Monosemy, like homonymy, has a number of problems associated with it. Perhaps most serious is that while it may well be that the distinct meanings associated with a particular form are related to a primary, abstract meaning, some meanings are demonstrably context independent. That is, although important, pragmatic knowledge alone is insufficient in predicting all of the distinct meanings associated with a particular form. For instance, in our examples for *over* it is difficult to see what kind of contextual knowledge would allow us to derive the spatial meaning of ‘above’, the non-spatial meaning of ‘again’ and the non-spatial meaning ‘finished’, all from a single, abstract meaning associated with *over*.

A second, and equally fatal, problem is that the primary meaning would need to be so abstract to be able to derive a set of such distinct meanings that it is difficult to see how the meanings associated with other spatial particles, such as *above* or *on*, could be mutually distinguished. Clearly, while accepting the insight that real-world pragmatic and contextual knowledge plays a significant role in the process of meaning construction, the linguistic evidence points to the conclusion that language users do store some distinct form-meaning pairings in long-term semantic memory. Hence, although the nature of meaning construction is a dynamic and highly creative process, not all meaning can be the result of situated (i.e., contextual) interpretation. Some of the interpretations associated with a particular form must be due to distinct meanings paired with the words themselves.

The position we advocate is that of polysemy. A polysemy approach suggests that the meanings associated with a spatial particle such as *over* are related in some fashion. We briefly preview our model below. The specifics are developed in chapter 3, and illustrated in detail by the analyses of English spatial particles throughout the course of this book. Our account of polysemy holds that a linguistic form is paired at the conceptual level, not with a single meaning, but rather with a network of distinct but related meanings. Hence, the meanings associated with a particular form constitute a semantic network. However, it is important to note that not all usages are contained within the semantic network. While some of the variation in uses of a word must be instantiated in long-term memory, and hence persist in the semantic network, some uses are created on-line in the course of regular interpretation of utterances. For instance, when we consider the semantic network for *over* in chapter 4, we will show in detail that the various senses in the text with which we began this chapter are instantiated in the semantic network, and hence are stored in long-term memory, whereas the meaning of moving from one side of an obstacle to the other in sentences of the following kind: *The cat jumped over the wall*, is a situated on-line interpretation, constructed for the purpose of local understanding in context.

We take the primary tasks for a model of the polysemy exhibited by spatial particles to constitute the following: (1) to establish what information is most appropriately included in the representation of the individual lexical entry and what information is appropriately represented as arising from cognitive processing and general world knowledge; (2) to model the systematic processes through which on-line contextually determined interpretations of spatial particles arise; (3) to model the systematic processes through which meaning is extended and through which the distinct senses – represented in long-term memory – become part of a lexical item's semantic network.

A major challenge for any theory of word meaning, and one we explore in detail in chapter 3, is to establish when a usage constitutes a distinct meaning component, which is legitimately instantiated in the semantic network, and

when a usage is simply a contextually derived interpretation constructed on-line. This is a methodological issue, but one which, as will be demonstrated, is tractable, in the light of linguistic evidence. As already intimated, another significant challenge is to establish what constitutes the primary sense for a particular semantic network. Again this is a methodological issue, and as we will see in chapter 3, we will rely on both linguistic and empirical evidence.

### **Interpretation of the utterance and the underspecification of meaning**

A cognitive approach to meaning construction holds that the interpretation of language is integrative, elaborative and inherently conceptual in nature. On this view, interpretation, which is to say meaning construction, is not simply the result of compositionally adding linguistic items. Rather utterances – lexical items and the syntactic configurations in which they occur – provide only minimal prompts for meaning construction. Language vastly underdetermines the rich interpretations normally assigned to even simple, de-contextualized sentences; sentential interpretation results from the integration and elaboration of these minimal linguistic cues at the conceptual level.

While the importance of pragmatic inferencing (e.g., implicature) and background knowledge in meaning construction is generally acknowledged, previous approaches to word meaning, both in the generative tradition (e.g., Pustejovsky, 1998) and in the cognitive linguistic tradition (e.g., Brugman and Lakoff, 1988; Lakoff, 1987), have failed to adequately take account of the largely non-linguistic nature of meaning construction, or what is more appropriately termed conceptual integration (Fauconnier and Turner, 1998, 2002). This has led previous scholars to fail to distinguish appropriately between information coded by the lexical item and information recruited from context, background knowledge and cognitive processing. As a result, these accounts have included a considerable amount of information in their representations of individual lexical entries, which is more appropriately understood as deriving from background and world knowledge, and human cognitive processing abilities (Kreitzer, 1997; Sandra, 1998; Sandra and Rice, 1995; Tyler and Evans, 2001b; Vandeloise, 1990). Sandra and Rice (1995), based on psycholinguistic experiments, have argued that such a degree of granularity is unwarranted, a view echoed forthrightly by Vandeloise (1990). As Kreitzer (1997) points out, the fine-grained distinction between instances of *over* as argued for in Lakoff (1987) provides a semantic network which is methodologically so unconstrained that ‘the model... [allows]... *across*, *through* and *above* all to be related to the polysemy network of *over*’ (1987: 292).

To fail to recognize that the source of much of the information which is necessary to establish an interpretation is not conventionalized information associated with a lexical item poses a significant problem. After all, if a theorist

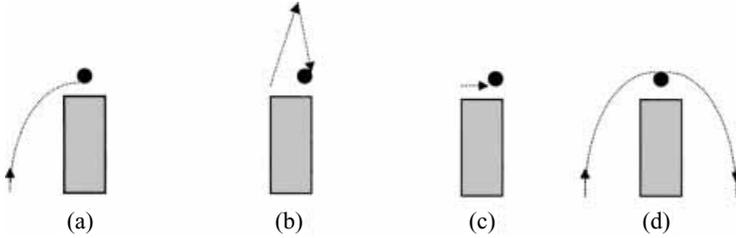


Figure 1.1 Possible trajectories for *The cat jumped over the wall*

believes that meaning is largely determined by language, then it follows that the theorist will attempt to explain meaning construction as deriving from the composition of lexical entries. But in order to produce the highly elaborate and complex interpretations that we regularly and ordinarily construct, the lexical entries would need to be spectacularly complex, as the next section will demonstrate. As we will see, assuming that lexical items are fully specified, that is, to assume that meaning is largely linguistic in nature, rather than conceptual, runs into immense difficulties, even in accounting for the correct interpretation of the most straightforward of sentences.

### A ‘simple’ example: the cat jumped over the wall

The following discussion demonstrates some of the problems encountered with an approach to sentence interpretation that relies on highly specified lexical entries. Consider the following sentence:

(1.1) The cat jumped over the wall.

In all probability, the reader will find this sentence unambiguous and readily understandable. Figure 1.1 presents four diagrams labelled (a) through (d). Before reading on, we ask that the reader select the diagram which best represents the event described by the sentence in (1.1).

We anticipate that the reader selected the fourth diagram (d). After all, the conventional reading of the sentence is that the cat begins a jump on one side of the wall, moves through an arc-like trajectory, and lands on the other side of the wall. Diagram (d) in figure 1.1 best schematizes this interpretation. On first inspection, this exercise seems straightforward, with no need for puzzlement or explanation. However, modelling how speakers of English consistently construct just the right interpretation, that is, (d), presents several complications. In the following, we will focus on how we know that the trajectory is the arc-shaped one, and not any of the others. In essence, where do we get this information from?

Even though the sentence in (1.1) is typically interpreted as unambiguous, it contains lexical items that have a range of interpretations. The behaviour

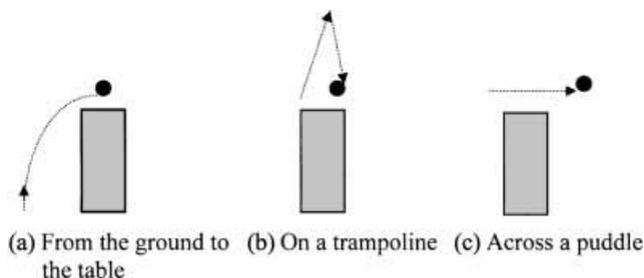


Figure 1.2 Trajectory paths potentially coded by *jump*

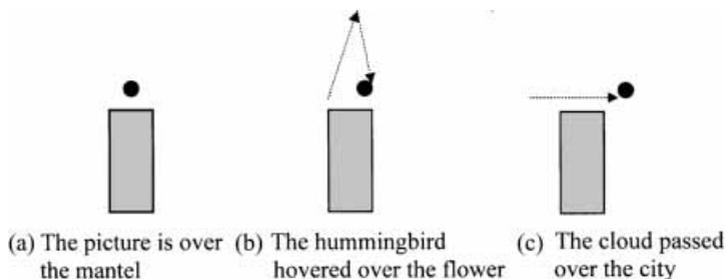


Figure 1.3 Trajectory paths potentially coded by *over*

described by *jump* has the potential to involve a variety of trajectory shapes, as in figure 1.2, some of which match the diagrams just rejected.

Similarly, figure 1.3 shows that *over* can be associated with several potential spatial configurations holding between the object in focus (i.e., *picture*, *hummingbird* and *cloud*) and the locating, backgrounded object (i.e., *mantel*, *flower* and *city*). Notice that these trajectory shapes also match some of the rejected diagrams from figure 1.1.

Thus, we face a seeming contradiction. The sentence in (1.1) which contains apparently ambiguous lexical items is consistently interpreted as unambiguous. The flip side of this contradiction is that, in spite of the range of different potential interpretations for *jump* and *over*, speakers of English consistently pick out just the right ones to assign interpretation (d) in figure 1.1 to the sentence *The cat jumped over the wall*.

Consider a further complication. Diagram (d) in figure 1.1 crucially represents the cat's motion ending at a point on the opposite side of the wall relative to the starting position of the jump. Yet, no element in the sentence explicitly provides us with this information.<sup>7</sup>

<sup>7</sup> Dewell (1994) argues that the primary mental representation associated with *over* involves the arced trajectory. He justifies this choice of primary meaning component in the following manner:

This exercise points to two fundamental and interrelated questions for an investigation of the role of lexical items in meaning construction. First, what is the source of information that is involved in the normal interpretation of sentences? And second, what constraints on interpretation must be posited in order to account for the native speaker's ability to choose consistently between competing interpretations of individual, polysemous items?

Many previous approaches to meaning construction have tended to assume what Jackendoff (1997) terms the strong version of the 'simple compositional' approach. This asserts that 'all elements of content in the meaning of a sentence' (1997: 48) are provided by the lexical items and the syntactic configuration in which they occur. In consequence, 'no aspects of a sentence's interpretation can arise from outside of the sentence itself' (1997: 41). Following this approach, let us consider just what kind of information it would be necessary to include in the lexical entries for *over*, *jump* and *cat* in order to obtain the interpretation diagrammed in (d) of figure 1.1.

First, we will consider *over*. Most previous studies of the polysemy exhibited by *over* assume that it codes the trajectory followed, as represented in a sentence like (1.1) (e.g., Brugman and Lakoff, 1988; Dewell, 1994; Kretzler, 1997;

'If we . . . simply imagine best examples of *over*, it seems that the "semicircular path" sense of *over* is the typical sense . . . It seems intuitively more accurate to posit a central schema that looks like [figure n.1]' (1994: 353).

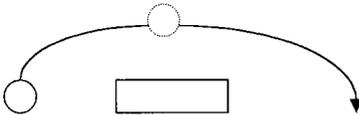


Figure n.1 After Dewell, 1994: 353

While we agree with much of the spirit of Dewell's analysis, in which information about the shape of the landmark (LM) – the entity which serves to locate the motile trajector (TR) in figure n.1 – and contact between the LM and the TR are eliminated, we are not convinced by his appeal to intuition. Our intuitions about the primary sense of *over* differ from Dewell's, as they do from Lakoff's and Brugman's. Clearly, intuitions concerning the primary sense of a spatial particle vary among analysts. Thus, Dewell's analysis fails to establish any criteria for determining the primary sense of spatial particles generally. Moreover, we find his primary mental representation problematic in that it actually contains two TRs – one at the beginning of the trajectory and one at the apex, as well as a LM. No explanation is offered for this dual TR representation. It seems that he is forced to this dual trajectory representation in order to capture a static Higher-than-Sense commonly associated with *over*, as well as a dynamic Above-across-Sense. Finally, while adjusting the primary representation of the trajectory for *over* from Lakoff's flat above-and-across trajectory to a semicircular trajectory allows improvements in the analysis of a subset of senses associated with the particular particle *over*, including a trajectory as part of the primary sense of a particle is problematic when applied to many other uses of *over*, as well as to many other spatial particles. This will become clearer as we analyse various spatial particles in later chapters. Our goal is to develop a more general, principled framework which can be applied to all English spatial particles. We feel that the tendency to focus on analysis of only one or two particles has led to explanations which might work with a particular item, but which tend to be ad hoc and not generalizable.

Lakoff, 1987). One reason why this may have been assumed is that a change of the spatial particle involved often results in a change in the interpretation of the trajectory, as attested in (1.2):

- (1.2) a. Jane marched up the stairs.  
b. Jane marched down the stairs.

In sentences such as these, information provided by the spatial particle clearly affects the interpretation of the trajectory. Theorists who assume that meaning must (predominantly) come from the sentential elements (i.e., from language) must also assume that it is the spatial particle which codes the trajectory, as with *over* in the example in (1.1). Hence, all permissible usages of *over* that exhibit configurational differences with respect to the shape of the trajectory and the landmark element (e.g., *the wall*) must be stored in memory as distinct senses. In such an approach, there is seemingly no room for non-linguistic information. To make this more concrete, let us consider some examples. Following Langacker (1987), we will refer to the focal element which follows the trajectory (e.g., *the cat*) as the trajector or TR and the backgrounded element as the landmark or LM. In one scenario, *over* can code a spatial relation in which the TR is located statically higher than the LM (as in *The picture is over the mantel*); in a second scenario, the TR is positioned higher than the LM while being in continuous motion (as in *The hummingbird hovered over the flower*); in a third, the TR moves on a trajectory which is above and across the LM (as in *The plane flew over the city*); in a fourth, in which there is contact between the TR and the LM, the trajectory is crucially shaped by the LM itself (as in *Sam crawled over the wall*), etc. However, such explanations run into difficulty, as they are unable to specify the shape of the trajectory in situations in which there is no contact between the TR and the LM (as is the case in sentence (1.1)).

Lakoff's (1987) fully specified account of *over*, for instance, assumed that the 'above-and-across' trajectory, as in sentences such as: *The plane flew over the city*, represented the primary sense of *over*. He argued that extensions of this 'above-and-across' trajectory should include information about the physical (i.e., metric) attributes and dimensions of the LM. Thus, different senses are suggested if the LM is vertical (e.g., a wall), extended (e.g., an ocean), vertical and extended (e.g., a city with high buildings), etc. Importantly, there is no claim that the trajectory shape is necessarily affected by the changes in the dimensions of the LM if there is no contact between the TR and the LM.

In a situation in which there is contact between the TR and LM, as in a sentence such as *Sam crawled over the wall*, the specifically mentioned LM provides the overall shape as well as the beginning and end points for the motion labelled as *crawling*. In order to specify an arc trajectory for 'above and across', when there is no contact between the TR and LM, additional semantic features (beyond those provided by the LM) which code where the movement started and

ended seem to be required. If such specification is necessary for one trajectory in which there is no contact, we might conclude that it should be included in the representation of all trajectories in which there is no contact. However, in the sentence *The plane flew over the city* the normal interpretation does not include information concerning where the motion associated with flying started and ended. If not all non-contact trajectories are coded by the spatial particle for beginning and end points, one might well wonder what criteria could be used in order to determine, in a non-arbitrary fashion, which uses of a spatial particle would need to have starting and end points of the trajectory specified. Even the most fully specified models of polysemy have shied away from this conundrum.

It might be possible to argue that the verb also carries information concerning trajectory shape. Previous accounts have been hesitant to do this. One reason for this is presumably because requiring the verb to code trajectory information would greatly expand the number of lexical listings or senses associated with each verb. In the case of *jump*, we have already seen three possible trajectories (as diagrammed in figure 1.2) and one can readily construct several more. In addition, there are instances in which the spatial particle crucially provides information about the direction of the trajectory (as in sentences (1.2a) and (1.2b)). Thus, placing information about the shape of the trajectory in the lexical entry for the verb would involve a good deal of redundancy, something most theories shy away from (Jackendoff, 1997).

Even if we set aside these difficulties and assume one could offer a principled account of *over* and *jump* which resulted in fully specifying all potential trajectory shapes, one is still faced with the question of how to account for selection of the correct sense within a particular sentence. A model which assumes the strong simple compositionality position might be able to account for the interpretation represented by (d) in figure 1.1 by somehow ‘coercing’ the appropriate choice given the other lexical items in the sentence.

One hypothesis for accomplishing this involves including information about the agent’s goals in the lexical entries for nouns (Pustejovsky, 1998). In the sentence in (1.1), beyond the basic four-legged-mammal-says-meow-type information, the entry for *cat* would need to include information about possible kinds of motion a cat could engage in. One subcategory of motion would include certain kinds of jumping. Since cats regularly engage in jumping, which entails varying trajectory shapes, information concerning the particular kind of jumping would have to be detailed. For instance, we might posit a putative taxonomy of the kinds of jumping and trajectories cats engage in. Consider some examples: Trajectory1: when the goal is to move from a lower position to a higher position (as from the floor to the table, say, as in diagram (a) in figure 1.2), the cat jumps in a roughly diagonal motion; Trajectory2: when there is no goal of forward motion (perhaps when startled), a cat can jump roughly straight up in the air (as in diagram (b) in figure 1.2); Trajectory3:

when the goal is to pass higher than and beyond a vertical impediment, the cat can jump with an arc trajectory (as in (d) in figure 1.1).

Nonetheless, even this amount of detail does not ultimately solve the problem because without knowing the cat's goal (which is not overtly signalled in the sentence in (1.1)), there is no way to rule out selecting the jump diagrammed in (b) in figure 1.1, for instance, in which the cat's trajectory begins on the wall, moves to a point relatively straight up and hence higher than the wall, and returns to the wall. That is, we would still be unable to predict that (d) in figure 1.1 is the normal interpretation assigned to the sentence in (1.1), namely *The cat jumped over the wall*.

A second attempt to coerce the correct selections might involve some kind of feature matching between *jump* and *over*. However, this solution runs into similar problems, as a match between *jump* in diagram (b) figure 1.2 and *over* in diagram (b) figure 1.3 cannot be excluded. Thus, we arrive at the inevitable conclusion that the information supplied by the syntactic configuration and individual lexical items, even when highly specified, cannot account for the interpretation normally assigned to this seemingly most straightforward of sentences.

### The role of background knowledge

Related problems in interpreting sentences in context led Grice (1975, 1978), Reddy (1979) and others to suggest that much of the normal interpretation of utterances does not derive from information coded by the utterance per se. Rather, they concluded that interpretation of ordinary sentences crucially involves humans drawing rational inferences based not only on what is uttered (the linguistic production), but additionally on the surrounding context, knowledge of speakers' intentions and knowledge of speakers' beliefs, including beliefs about how the world works (see especially, Reddy, 1979). As Green (1989) notes, virtually all natural language utterances are vague and ambiguous. Speakers must always add information to the linguistic elements present in an utterance in order to establish an appropriate interpretation. Grice has articulated this general approach to natural language interpretation in terms of the Cooperative Principle and the maxims which he saw as particular instances of the Cooperative Principle. More recently, Sperber and Wilson (1986) have argued that Grice's insights are more appropriately framed in terms of the single principle of relevance. (See Green, 1989 and Sperber and Wilson, 1986 for a full discussion.)

A simple compositional approach to meaning construction and lexical items advocates including in the lexical entry all information a speaker would require in order to establish the appropriate interpretation of any sentence in which the lexical item occurs. This position forces inclusion of vast amounts of information in the mental lexicon. Setting aside the cumbersome nature of such

a model, even then, the listener would still be required to make inferences about the speaker's intentions and beliefs, and the relevance or probability of a particular interpretation within the exact context in which the utterance is issued.

As we have just seen, even normal interpretation of simple, de-contextualized sentences in which spatial particles occur seems to involve information that is not explicitly provided by the individual lexical items. Although our work departs from the theorists cited above (e.g., Grice and those who accept his position that pragmatics supplements a truth-conditional semantic component) in many important ways, we take their conclusion that linguistic utterances radically underdetermine the meaning involved in normal interpretation of utterances as both foundational and fundamental.

Some of the strongest support for this general position has come from the field of experimental psychology. Starting in the early 1970s a number of experimental psychologists such as Frank and Bransford (1973), Rummelhart (1975), and Wilson and Anderson (1986) demonstrated the importance of background knowledge and expectation in interpreting connected text. For instance, numerous experiments established that otherwise vague text could readily be assigned a reasonable interpretation if the reader were provided with a relevant title, such as 'Doing the Laundry'. Lacking this background frame, readers with comparable reading skills found the same text confusing and difficult to interpret. In other experiments, readers with background knowledge relating to a text were consistently shown to remember more information and make more appropriate inferences than readers with comparable reading skills who lacked the appropriate background knowledge. Many researchers have come to the conclusion that the interpretation of text represents a synthesis of knowledge in the reader's/speaker's mind and the information provided by the linguistic code.

We hypothesize that, for a sentence such as *The cat jumped over the wall*, the nature of the normal interpretation constitutes meaning construction conducted at the conceptual level. Moreover, we will argue that the processes which mediate and facilitate such a conceptualization – we equate meaning with conceptualization – are conceptual (rather than linguistic) in nature.

This position stands in sharp contrast to theories which argue that interpretation of sentences relies primarily on the cumulative information supplied by the individual lexical items and the syntactic configurations in which they occur (compositional semantics). We will argue that the linguistic utterance acts as a minimal prompt for conceptual construction which is far richer than the combined information provided by the lexical items.

Normal interpretation of even a simple sentence, such as *The cat jumped over the wall*, is crucially tied to basic, recurring experiences with the world. A major part of this experience involves understanding force dynamics, such as

gravity, and how these dynamics affect physical objects, such as cats (Talmy, 1988a, 2000). Given our recurring experiences with gravity and how objects move, we know that the action we label as *jump* involves pushing off from a surface, momentarily leaving the surface, and eventually returning to a stable surface. A fundamental component of the semantics of *jump* is that the TR is physically displaced, that is, motion is involved, and hence a trajectory is projected. The linguistic prompt *over* in sentence (1.1) provides the specific, key information that, at some point in the trajectory, the cat is higher than the wall. Our knowledge of gravity and cats tells us that a cat, unlike a hummingbird, cannot suspend itself in midair for long periods of time. The process of conceptual integration of the information prompted for by *over* with our knowledge of objects and force dynamics results in the conceptualization represented by our diagram (d) in figure 1.1. Thus, we argue that none of the individual lexical items explicitly provides information concerning the shape of the trajectory. Rather, we will argue that this information, and indeed a good deal of the information needed to establish normal interpretation of most sentences, comes from cognitive processes, conceptual structure and background knowledge rather than the individual lexical items. We will return to *The cat jumped over the wall* in chapter 4.

Throughout this book, we will argue that attempting to list detailed information in the basic lexical representation of spatial particles fails to account for everyday meaning construction. In contrast, we posit lexical representations which are more abstract in nature. Our analysis models detailed knowledge of the spatio-physical world which forms part of the normal interpretation of utterances via application of a set of inferencing strategies and ways of constructing (i.e., seeing) *spatial scenes*. Spatial scenes, such as the scene prompted for by a sentence like: *The cat jumped over the wall*, involve conceptualizing a spatio-configurational relation between entities we encounter in the world around us and with which we interact. Hence, a spatial scene is a conceptualization grounded in spatio-physical experience. This analysis allows us to avoid the problems encountered with a more fully specified lexical representation while revealing the systematic semantic relations among the many meanings of individual English spatial particles.

Nevertheless, we must emphasize that we are not so much relegating the role of lexical items in meaning construction to an unimportant place, as assigning them their appropriate place. To say that lexical items act as prompts for meaning construction is not to say that lexical forms do not crucially contribute to the meaning-construction process. Clearly, lexical items, in general, and spatial particles, in particular, do contribute meaning. For instance, *The cat jumped over the wall* is regularly assigned a different interpretation from *The cat jumped beside the wall*. However, we must also be aware that what they contribute is a prompt for a complex conceptual elaboration.

### Dictionaries versus encyclopedias

The distinction between treating lexical items as fully specified versus seeing them as merely prompts for complex conceptual elaboration has been framed by some scholars in terms of a distinction between a dictionary versus an encyclopedic view of word meaning (cf. Haiman, 1980; Langacker, 1987; Wierzbicka, 1988). On this view, linguists who subscribe to the dictionary view of word meaning attempt to identify a restricted and finite set of specifications that constitute the linguistic knowledge properly associated with the lexical entry for a particular lexical item. However, as both Haiman (1980) and Langacker (1987) observe, attempts to restrict and, hence, determine which specifications should be included in a particular lexical entry and which should be excluded is impossible on practical grounds. Moreover, as the foregoing discussion has highlighted, such attempts will inevitably fail to account for the variety and range of distinct interpretations ordinarily associated with a particular lexical item.

The mistake in adopting a dictionary view of lexical items has been to view words as ‘containing’ meaning, a naïve view of communication, which Reddy (1979) argued was to fall prey to what he termed the conduit metaphor. We suggest not that words contain meaning, but rather, in the spirit of scholars such as Fauconnier (1997), Langacker (1987) and Turner (1991), that words *prompt for* highly complex conceptualizations. As Langacker has felicitously put it: ‘linguistic expressions are not meaningful in and of themselves, but only through the access they afford to different stores of knowledge that allow us to make sense of them’ (1987: 155). This view, which might be termed the encyclopedic view of word meaning, treats lexical items as *points of access* (in Langacker’s terms) to the totality of our knowledge regarding a particular conceptual entity. This reflects what we know about how a particular linguistic expression is used and our knowledge of that aspect of the conceptualized world which the entity it prompts for inhabits. Moreover, such knowledge is accessed in conjunction with various inferencing strategies which allow us to build elaborate conceptualizations in ways maximally coherent with, and contingent upon, our experiences of the world.

One way in which we will attempt to capture the encyclopedic knowledge prompted for by a particular spatial particle is to model the meaning component associated with a form such as *over*, for instance, in terms of a semantic network (we will develop the notion of a semantic network in detail in chapter 2). That is, a lexical item should be thought of as prompting for a range of meanings, the particular meaning selected being determined by conceptual integration in context. Moreover, in chapter 3 we will argue that within a semantic network not all of the meanings associated with a particular spatial particle have equal status. For instance, we will suggest that some meanings can be determined as

being what we will term more *primary* than others (the notion of *primariness* will be taken up in chapter 3).

In essence, then, in normal communication lexical items do not occur in isolation. In point of fact, when humans use lexical items, the lexical items always occur in context and their precise interpretation changes with each use. Hence, a typical dictionary definition inevitably fails to provide for the infinite amount of variation and detail that arises when a lexical item is interpreted in context. Put another way, a lexical item prompts for a highly specified conceptualization. Crucially, however, this rich specification emerges at the conceptual level, due to integration of prompts in context, and is vastly more complex than anything which can be derived or predicted from the individual lexical items integrated in compositional fashion. Accordingly, it is to the notion of conceptualization which we now turn.

### **The conceptual nature of meaning**

Within the study of linguistic semantics, both in the philosophical and linguistic traditions, it has been widely assumed that meaning derives from the fact that language refers directly to the world. The means whereby language ‘matches up’ with the world has relied on the notion of truth. Yet, two fundamental problems fatally undermine this approach. In terms of the study of meaning, some scholars have suggested that so-called truth-conditional or model-theoretic approaches to semantics are concerned with what has been termed informational significance rather than cognitive significance, and hence do not represent the study of meaning, properly conceived. As Wierzbicka (1996) puts it, ‘truth-conditional’ semantics ‘doesn’t seek to reveal and describe the meanings encoded in natural language, or to compare meanings across languages and cultures. Rather, it sees its goal as that of translating certain carefully selected types of sentences into a logical calculus. It is interested not in meaning (in the sense of conceptual structures encoded in language) but in the logical properties of sentences such as entailment, contradiction, or logical equivalence’ (1996: 8).

A second problem is that a truth-conditional approach assumes that much of language directly reflects and refers to the world. As the cognitive scientist Gilles Fauconnier observes: ‘When language expressions reflect objective events and situations, as they often do (and often do not), they do not reflect them directly, but rather through elaborate human cognitive constructions and construals’ (1997: 8). This point has been elaborated in detail by Ray Jackendoff (1983, 1990, 1992). Jackendoff has pointed out that one of the most important insights to emerge from the work on perception is that our perceptions of the world are determined largely by conceptual organization being imposed on sense-perceptory input. That is, what we directly experience is not an objectively real world. Rather, what we experience as everyday reality is mediated and shaped

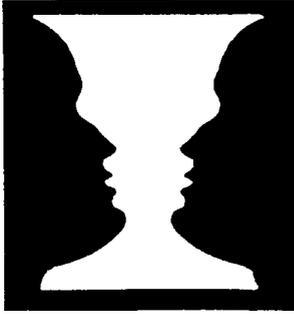


Figure 1.4

by human conceptual organization to which we necessarily and unconsciously subject sense-perceptory input (cf. Dennett, 1991; Putnam, 1981).<sup>8</sup>

In essence, the patterns and organization we perceive as reality do not in fact exist independently in the world itself, but are largely the result of our cognitive processing. For instance, in figure 1.4, the reader might see either a vase or alternatively two faces. Yet, the particular image seen is not dependent upon the raw input, which presumably remains unchanged. What changes is the organization of our perceptions, by mental operations to which we do not have conscious access. These present us with two alternating and conflicting interpretations of experience. It would be erroneous to claim that the vase interpretation is true and the two-face interpretation is false, or vice versa. It is also contradictory to say both exist simultaneously; by fixating on the figure what is seen will alternate between the two perceptions, but we cannot perceive both simultaneously. Clearly, figure 1.4 is not a drawing of something that exists in the world. It is we who perceive it to be of something. This is instructive as it demonstrates that although there is a world of sense-perceptory information out there, what we in fact perceive is determined by how we unconsciously organize and hence make sense of the input.

<sup>8</sup> In other words, as humans we only have access to our conceptual system. Words (linguistic elements) reference concepts. Concepts are not, however, unrelated to the 'world out there'. We believe that concepts are best understood as arising from redescribed percepts. Many percepts arise from sensorimotor experiences derived from the world. Because of humans' particular physical and neurological architecture, we perceive objects and actions in particular ways. Percepts can also arise from internal states, such as an emotional state.

These percepts are the raw data which, when reanalysed, form the concepts to which we have direct access. When we 'refer' to some object or event which we have perceived in the 'world out there', we are 'referring' to a mediated percept (i.e., one which has been filtered through our particular human neurological apparatus) which has, in turn, been redescribed into a format accessible to our conceptual system. It is only once the raw stimulus from the 'outside world' has been so mediated and redescribed that it can be assigned a linguistic label, such as *dog*, for instance.

This insight has profound consequences for a theory of meaning and language. If our world of experience is not the real world itself, but the real world as mediated by our cognitive faculties, then the world to which we have direct and conscious access is the mental world of experience, which is to say the conceptual system. Jackendoff terms this our projected world. When we use language, then, we are referring to concepts in our projected world, which indirectly reflects the real world. Jackendoff has felicitously summarized this position as follows:

[W]e must take issue with the naïve [truth-conditional] position that the information conveyed by language is about the real world. We have conscious access only to the projected world – the world as unconsciously organized by the mind; and we can talk about things only insofar as they have achieved mental representation through these processes of organization. Hence the information conveyed by language must be about the projected world. We must explain the naïve position as a consequence of our being constituted to treat the projected world as reality. (Jackendoff, 1997: 29)

If language cannot refer to an objective world, precisely because we have no direct access to such, then language prompts for concepts. Moreover, linguistic elements, as we have noted, consist of form-meaning pairings (where the meaning component constitutes a semantic network). That is, words (and constructions more generally) are symbolic assemblies consisting of a phonological pole and a semantic pole (Langacker, 1987, 1991b).<sup>9</sup> In line with the assumptions set forth at the beginning of the chapter, the semantic pole derives from conceptual structure. We conclude from this that the semantic value of a lexical item can be equated with a particular concept. This conclusion has now been reached by an increasing number of scholars who have recognized the fundamentally conceptual nature of language and conceptual representation (e.g., Heine, 1997; Fauconnier, 1997; Jackendoff, 1983, 1990; Langacker, 1987, 1991b; Lakoff, 1987; Talmy, 2000). The cognitive linguist Ronald Langacker (1991a) has summarized the position as follows: ‘Semantic structures [meanings] are conceptual structures established by linguistic convention – the form which thoughts must assume for purposes of ready linguistic symbolization. Thus, semantic structure is conventionalized conceptual structure’ (1991: 108–9). In other words, lexical items prompt for conventionalized concepts.

In order to make the claim that words prompt for concepts more concrete, let us consider an example. Consider the word *bird*. The concept which corresponds to this linguistic form is but a sketch. The details of shape, size, vocalization ability, even ability to fly are filled in by contextual and real-world

<sup>9</sup> Langacker argues that not only words but any kind of complex expression, such as certain grammatical constructions, are symbolic assemblies. By saying that words are symbolic assemblies consisting of a phonological pole and a semantic pole, we are not claiming that they are the only such symbolic assemblies.

knowledge. Semantic representation (word meaning) provides only a skeletal prompt, which subsumes little more than the scaffolding for the construction of meaning. Lexical items prompt for conceptualization, the process whereby rich and elaborate meanings are constructed.

The assumption that semantic representation prompts for a conceptually mediated representation of the world also provides powerful insights into, and accounts for, many uses of spatial particles which have previously been labelled as arbitrary. If one assumes that language directly reflects the real world, then one assumes that the objectively metric properties and principles of Euclidean geometry which appear to hold for the spatio-physical world ‘out there’ will form the basis of linguistic descriptions of spatial scenes and uses of spatial particles. Talmy (1988b, 2000) has argued persuasively that conceptualized space as reflected in language is not Euclidean in nature, that is, it is not held to notions of fixed distance, amount, size, contour, angle, etc. He argues that conceptualized space is topological in nature, that is, conceptualized space ‘involves relativistic relationships rather than absolutely fixed quantities’ (1988b: 170). Assuming that language refers to conceptual structure provides the insight that the relationships between objects are subjective and largely influenced by the interpretation imposed by the conceptual system.

A further benefit of assuming that language is conceptual in nature is that we now have a means of distinguishing between mundane, yet ubiquitous sentences such as: *Jane stood in the flower-bed*, versus: *Jane stood on the flower-bed*. If we assume that there is a direct relation between the real world and language, as in truth-conditional approaches, there is no explanation for why English speakers can describe the event of a person standing such that her feet are in contact with the piece of ground designated as the flower-bed, using either *on* or *in*. Traditional approaches have assumed that examples such as these are semantically equivalent. However, we are now able to see that each sentence represents a distinct conceptualization (or construal in the sense of Langacker, e.g., 1987) of an objectively identical scenario. This is analogous to the way in which in figure 1.4 we were able to see either the vase or the faces. What we see is mediated by the conceptual system, which has a number of ways to represent the same scene. These issues will be pursued further in chapter 3.

### Conclusion

We began this chapter by demonstrating that spatial particles typically have numerous meanings associated with them. We claimed that a subset of interpretations represent those meanings which must be stored in memory, and hence are permanently available. These meanings we termed senses. We also suggested that some meanings associated with words must be due to pragmatic inferencing, context and background knowledge. These meanings, we

suggested, are constructed on-line in the moment of speaking and listening. In attempting to distinguish between the meaning contributed by language and the meaning due to world knowledge and cognitive processing, we saw, as with our illustration of *The cat jumped over the wall*, that many previous accounts have vastly underestimated the amount of information which is not accounted for by the conventionalized meaning of lexical items and the grammatical construction in which the lexical items occur. This led us to the general conclusion that meaning construction must be inherently conceptual in nature. This also points to a finding which is coherent with studies which are broadly ‘cognitive’ in the sense adduced – namely that meaning is fundamentally mental in nature, referencing conceptual structures rather than directly referencing entities inhering in an objectively verifiable and mind-independent world. Language refers to conceptual structure, which indirectly reflects the world. (See Evans, 2000: chapter 2; Jackendoff, 1992: chapter 12.) These general findings have profound consequences for a theory of word meaning, a theory of meaning construction, and perhaps most crucially, for an understanding of the relationship between language, thought and the nature of reality.