

Semantic Relations and the Lexicon

Antonymy, Synonymy, and Other Paradigms

M. Lynne Murphy

University of Sussex



published by the press syndicate of the university of cambridge
The Pitt Building, Trumpington Street, Cambridge, 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>

© M. Lynne Murphy 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^AT_EX 2_ε [tb]

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

Library of Congress Cataloging-in-Publication data

Murphy, M. Lynne.

Semantic relations and the lexicon: antonymy, synonymy, and other paradigms /

M. Lynne Murphy.

p. cm.

Includes bibliographical references and index.

ISBN 0 521 78067 5

1. Semantics. 2. Lexicology—Psychological aspect. I. Title

P325.M826 2003 401'.43 – dc21 2003046079

ISBN 0 521 78067 5 hardback

Contents

<i>Acknowledgments</i>	page viii
<i>Symbols and typographical conventions</i>	ix
Part I Paradigmatic relations, generally	
1. Why lexical relations?	3
2. A pragmatic approach to semantic relations	25
3. Other approaches	61
Part II Paradigmatic relations, specifically	
4. Synonymy and similarity	133
5. Antonymy and contrast	169
6. Hyponymy, meronymy, and other relations	216
7. Lexicon and metalexicon: implications and explorations	237
Appendix: Relation elements	243
<i>Notes</i>	245
<i>References</i>	254
<i>Index</i>	274

1 Why lexical relations?

“How are you doing?” I would ask.

“Ask me how I am feeling?” he answered.

“Okay, how are you feeling?” [. . .]

“I am very happy and very sad.”

“How can you be both at the same time?” I asked in all seriousness, a girl of nine or ten.

“Because both require each others’ company. They live in the same house. Didn’t you know?”

Terry Tempest Williams, “The village watchman” (1994)

As for any other phenomenon in the world, the existence of paradigmatic semantic relations among words calls for some kind of explanation – or perhaps several kinds of explanation. Are these relations among words, or among the things the words represent? Are the relations arbitrary or rule based? Language specific or universal? A product of linguistic or general cognition? These questions are the focus of this book. First, however, we must ask what these questions mean, and why we might care to trouble ourselves with them.

As linguistic theories have progressed in modeling human language ability, the lexicon has become more central to those theories. With this new or renewed attention to the mental lexicon, two problems become evident. Firstly, there is no generally accepted theory of how the lexicon is internally structured and how lexical information is represented in it. Secondly, the lexicon must interface with the conceptual system, but there is little agreement about which information should be included on which side of the lexical-conceptual boundary, how conceptual information is represented, and even whether a lexical-conceptual boundary exists.

At the very least, most interested parties agree that the paradigmatic semantic relations among words – antonymy, synonymy, hyponymy and the like – are somehow relevant to the structure of lexical or conceptual information. Beyond this vague statement of “relevance,” however, opinions, assumptions, and models vary drastically. For some investigators (e.g., Katz 1972, Kempson 1977, Pustejovsky 1995) accounting for such relations is one of the purposes of lexical semantics, just as accounting for relations like entailment and contradiction is a

foundational problem in sentential or propositional semantics. For others (e.g., Deese 1965, Lehrer 1974, Mel'čuk 1996, Fellbaum 1998c) relations among words constrain or determine meaning, rather than vice versa. These positions are often stated as background to other discussions, as if they are uncontroversial. However, the differences between them underscore the extent to which the genesis, representation, and uses of paradigmatic relations are as yet unsettled matters for linguistic and psycholinguistic theory.

The following chapters have three purposes: (a) to bring into focus the various theoretical positions on paradigmatic semantic relations, (b) to summarize and analyze research about them from a range of disciplines and methodologies, and (c) to present a new, pragmatic approach to these relations. In this chapter, I examine the implications of taking a pragmatic and psycholinguistic perspective on semantic relations, define some of the vocabulary used here, and justify some assumptions about the mental lexicon and the conceptual system. The final section outlines the remainder of the book.

1.1 Approaching semantic relations

Semantic relations among words have captured the interest of various brands of philosophers, cognitive psychologists, linguists, early childhood and second language educators, computer scientists, literary theorists, cognitive neuroscientists, psychoanalysts – investigators from just about any field whose interests involve words, meaning, or the mind. The good news, then, is that we can access a broad and detailed literature that approaches the topic from a variety of methodological and theoretical perspectives. The bad news is that each of these perspectives carries its own implicit assumptions about why semantic relations are interesting, how they are (or are not) relevant to the structure of language or thought, and what research methodologies are (and are not) valid or revealing. So, while I report research from several of these fields, it is important to define the particular perspective taken here before discussing the literature or presenting new hypotheses. Doing so not only makes the presentation more comprehensible, but also serves as an acknowledgment that examining work by others often entails reading it from a different perspective from that in which it was written. In the following chapters, research that originated in a variety of fields and perspectives is critically assessed in light of the assumptions introduced in this chapter.

The overarching goal here is to provide an account of how individuals know (or determine) whether words are semantically related or not and, if they are related, what type of relation is involved. In other words, on what bases are judgments of semantic relatedness made? The perspective taken is pragmatic and psycholinguistic.¹ By **psycholinguistic**, I mean that the goal is to provide a psychologically plausible model of the knowledge and processes involved

in semantic relations phenomena in human language behavior. So while structuralist, formal, and computational models of the lexicon are among those considered here, they are assessed here on the basis of whether they model human language abilities in a mind-like way. By **pragmatic**, I mean that the linguistic phenomena described here are considered with reference to their use and their status in a human mind within a human culture. This contrasts with strictly formal or descriptive accounts of semantic relations, in which words are considered only with reference to their definitional meanings and those definitions' relations with each other. Thus it will not be sufficient here to say *hot* is the antonym of *cold* because our models of the words' semantic properties fulfill the formal requirements for antonyms. As discussed in section 2.1, semantic relations among words can depend on more than just the semantic qualities of a word, and they are highly context dependent. So we must discern (for example) how one determines in a certain context that *hot* is the best candidate to be the antonym of *cold*, but in another context *warm* or *cool* or something else might be a better opposite for *cold*. In essence, this means that I do not start from the position of considering semantic relations as a matter of analytic or objective truth, but instead as a matter of language users' idiosyncratic mental representations (and processes involving them), which can be investigated through their judgments and behavior. While paradigmatic semantic relations have been defined in logical terms (with varying degrees of success – see Part II), such definitions reveal little about the roles of semantic relations in lexical memory and language use.

The pragmatic and psycholinguistic perspective, then, is concerned with the relationships between competence and performance. Studying these relationships involves determining *what* one must know in order to know *how* to do something (like produce or interpret a meaningful utterance) and *what* we know as a result of having done this thing. The English language, unfortunately, is not very helpful in making plain the differences among these (and other) kinds of knowing. For the following discussion, at least four kinds of knowledge are relevant. **Fixed mental representations** in long-term memory are needed for some types of knowledge of language. For instance, for any word in my active vocabulary, I must have some representation of its phonemic structure in long-term memory.² For example, I know that *night* is basically pronounced [najt] because I have some fixed mental representation of this fact of English. Knowledge of language also involves **procedural knowledge**, which linguists usually represent as rules. So, for example, I know that most English plurals are made with *-s*, and I know to vary the pronunciation of the plural marker in accordance with the morphological context. These first two types of knowledge allow for a third kind: **generated mental representations**. So, once I use my ability to make *night* plural, I have a mental representation of this plural in my short-term memory (which may come to be stored in long-term memory as

well). All of the foregoing types of knowledge do not necessarily involve the fourth type: consciousness or **awareness** of the representations or processes involved. Of course, if we were aware of these rules and representations, we would not need to do much linguistic research, since the answers to our research questions would be plainly evident. In awareness, some “higher-level” part of the mind has access to some “lower-level” subconscious part. For instance, my awareness that a book is in front of me is dependent on all sorts of knowledge that I am not aware of, including perceptual processes and representations and their interface with my conceptual representation of what a book looks like. Awareness is the least interesting type of knowledge for our current purposes, since it is not so much about what we know, but what we know we know (and knowing about knowing is a problem for epistemologists, not linguists). If we are aware that *night* is the antonym of *day*, it is because the conscious mind has some access to what is going on in the subconscious mind.

Knowing that two words are antonyms or synonyms could involve any of the subconscious types of knowledge. If such relations are mentally fixed, then we either know them because we were innately programmed with this knowledge or because we learned that the two words are related and added that information to our mental representations of these words. We can rule out innate representation of lexical relations, since the knowledge is language specific. *Hot* and *cold* cannot be innately programmed as antonyms, since this fact is only relevant to English speakers. Having innate mental representation of every relation for every possible language is plainly impossible since there is an infinite number of possible languages. Even if we suppose that only semantic information (not words per se) is opposed in the semantic relations, the knowledge is still too language specific, since the particular semantics of *hot* are quite different from the particular semantics of French *chaud* (see chapter 5 and Cruse 1986) or Chinese *rè* (Prator 1963).³

This leaves us with two possibilities as to how we know that two words are semantically related. We may know the relation because we learned it as fact, just as we learn other facts about words such as their pronunciation or part of speech. In this case, experience of the words in relation is recorded in long-term memory. So, for instance, I might know that *hot* and *cold* are antonyms because I heard them being used in contrast and I (subconsciously) made this information part of my mental representation of these words. Another possibility is that semantic relations among words are generated. The knowledge that two words are antonyms would then involve a generated mental representation based on some set of rules or principles for generating relations among words. In this case, my knowledge that *hot* and *cold* are antonyms is something that I re-generate whenever the need arises. As discussed in chapter 2, neither of these possibilities alone is sufficient to explain our linguistic performance with respect to semantic relations. While I argue for principles that generate instances of

semantic relations, the derivability of relations does not preclude experiential learning and fixed mental representation of some relations as well. So, while all semantic relations among words can be generated via a single relational principle (introduced in chapter 2), this does not mean that the relations have to be generated each time they are needed.

A pragmatic and psycholinguistic perspective on semantic relations entails an interdisciplinary approach, since this perspective acknowledges that language must be considered with reference to social, communicative, and psychological constraints. Thus a wide range of types of evidence is available for determining how semantic relations are mentally represented and linguistically expressed. Among the types of evidence to be discussed are:

- Speakers' judgments of semantic relatedness
- Corpus-based studies of semantically related words
- Descriptions of semantic relations in thesauri and dictionaries
- Tests of computational models of lexical knowledge
- Psycholinguistic experimentation that is intended to reveal lexical organization (e.g., word association, lexical priming)
- Naturally occurring and experimental data on language acquisition
- Naturally occurring and experimental data on lexical loss or speech errors caused by anomia, aphasia, and run-of-the-mill disfluency
- Discourse analyses of the uses of semantic relations.

Each of the above sources of information has its own limitations. Speakers' intuitions and judgments are notoriously at odds with real language use (Labov 1975). Corpus studies often over-rely on written sources and tend to assume that the form of a lexical item is graphic, rather than phonemic. Dictionaries and thesauri reflect conscious reasoning about language and usually have commercial and practical missions that conflict with their descriptive usefulness. Computational, psychological, and neurological studies frequently fail to distinguish between linguistic and non-linguistic knowledge, since they often use words to represent concepts. One challenge here is to identify which work from other perspectives can be integrated into a discussion in the pragmatic perspective. That is, we must be cautious regarding interdisciplinary studies, and even intradisciplinary ones, because some alleged evidence for or against a position might be uninterpretable or irrelevant when considered in the framework of the current discussion. Another challenge is to use information from fields other than linguistics with caution and humility. As noted by Pederson and Nuyts, "There has been an increased sharing of methodological techniques across the traditional disciplinary boundaries . . . However, such techniques are all too often borrowed without a clear sense of their strengths, weaknesses, and underlying theoretical assumptions" (1997: 6). As is clear throughout the following chapters, my intellectual biases are toward linguistics, and my foremost purpose is to contribute to that field of inquiry. Nevertheless, I hope that the

cross-disciplinary nature of the material discussed will make this work useful to readers in any of the cognitive sciences.

1.2 Relations and 'nyms: some definitions

So far, the topic of study has been described as *paradigmatic semantic relations among words*. In the literature, these relations are usually called *lexical relations* or *semantic relations*, and sometimes those two terms are used contrastively. The common element, **relation**, is fairly vague, but in its most basic use it describes co-membership in a definable set. So, for example, *sky* and *high* are related in that they are members of the set of English words that rhyme with *eye*. *Relation* is also used to distinguish the types of definitional criteria that define such a set. So, the relation between *sky*, *high*, and *eye* is the rhyme relation (i.e., the criterion for membership in the relational set is similarity of word-final sounds). For our purposes, *relation* can stand for **paradigmatic relation**, in which the set of words forms some sort of paradigm, such as a semantic paradigm that contains members of the same grammatical category that share some semantic characteristics in common, but fail to share others. So, for example, the set of basic color terms forms a paradigm whose members are adjectives (or nouns), each referring to a different section of the color spectrum. Not all paradigms are semantically defined, of course. Inflectional paradigms, for instance, include the possible variations of a lexical item in some inflectional category, such as number. So a morphological paradigmatic relation exists between *child* and *children*. Paradigmatically related words are, to some degree, grammatically substitutable for each other. For example, *blue*, *black*, and any other member of the color paradigm can sensibly and grammatically occur in the phrase *a — chair*. In this way, paradigmatic relations stand in contrast to **syntagmatic relations**, which are relations between words that go together in a syntactic structure. For example, we can speak of a syntagmatic relation between *eat* and *dinner*. The two types of relation are not always easy to distinguish (see 2.1.5), although the (debatable) rule of thumb for distinguishing them is that paradigmatic relations hold between members of the same grammatical category, while syntagmatic relations involve members of different grammatical categories.

For present purposes, it makes sense to use the term **semantic relations** to indicate relations defined by semantic paradigms – but not before issuing some caveats. *Semantic relations* is sometimes used to denote phrasal or sentential relations such as paraphrase, entailment, and contradiction, but here it should be understood to mean ‘paradigmatic semantic relations among words.’ Given the pragmatic perspective taken here and the fact that non-semantic factors may affect these so-called semantic relations (see 2.1.5), one might argue that they should be called *pragmatic relations*. But that term misses the point that even

if non-semantic factors (such as phonetic form or register) come into play in antonymy or synonymy, the most basic requirement is semantic relatedness. Non-semantic factors may affect judgments of how well a set of, say, synonymous words exemplifies the synonym relation, but the meanings of the words make or break the relation.

The term **lexical relation** is used here to indicate any paradigmatic relation among words, not just a semantic relation. So, lexical relations include phonetic relations (such as rhyme or alliteration), morphological relations (such as inflectional variation), and morpho-syntactic relations (such as co-membership in a grammatical category). Again, a caveat is in order. The term *lexical relations* is ambiguous, in that it could refer to relations among words (on a page, in a mind, or wherever they might exist) or to relations (among lexical items) within the mental lexicon. For some authors, the two meanings are interchangeable, since they hold (or assume) that if words are related, then that relation is represented in the lexicon (see 3.3). However, I (in chapter 2) take the position that relations among words are not among the types of information about words that can be represented in the lexicon. This position contrasts with that of, for example, Gross, Fischer, and Miller (1989). They distinguish between antonym pairs like *big/little* and other semantically opposed pairs, such as *gigantic/tiny*, claiming that the former are lexical antonyms (i.e., intralexically related) as well as conceptual opposites (semantically related), while the latter are only conceptually opposed. For them, this means that the *big/little* contrast must be represented in the mental lexicon, but the relation between *gigantic* and *tiny* is not a part of those words' representation in the lexicon. In the context of the term *lexical relations* in this book, **lexical** should only be assumed to mean 'involving words' rather than 'contained in the mental lexicon.' The term **intralexical** indicates that a structure or piece of lexical information is contained within the lexicon. **Metalexical** indicates information that is not contained in the lexicon, even though it may be information about words.

The main relations discussed here are exemplified as follows:

synonymy: *sofa=couch=divan=davenport*

antonymy: *good/bad, life/death, come/go*

contrast: *sweet/sour/bitter/salty, solid/liquid/gas*

hyponymy, or class inclusion: *cat < mammal < animal*

meronymy, or the part-whole relation: *line < stanza < poem*

The equals sign (=) is used to indicate **synonymy**. The slash (/) between members of antonym or contrast sets signifies the semantic incompatibility of the contrasting words. **Antonymy** is a subtype of **contrast**, in that it is contrast within a **binary** paradigm. While the term *antonymy* is sometimes reserved for more specific relations, it is used here for any binary semantic contrast among lexical items (whereas *opposite* is used more broadly here, not limited to contrast between linguistic expressions – see 2.2.2). The 'less than' sign (<)

in the hyponymy and meronymy examples indicates that these relations are hierarchical and asymmetrical. That is, *stanza* is a meronym of *poem*, but *poem* is not a meronym of *stanza*. The converse relations of hyperonymy and holonymy can be represented by the 'more than' sign (>), as a *poem*>*stanza* (i.e., '*poem* is the holonym of *stanza*'). For example, *cat* does not have the same relation to *mammal* (*cat*<*mammal*) as *mammal* has to *cat* (*mammal*>*cat*). In one direction, it is a relation between a category and its superordinate category, and in the other, it is a relation between a category and its subordinate. On the other hand, synonymy, antonymy, and contrast are **non-hierarchical** relations, and are usually characterized as **symmetric** relations in that the relation between, say, *couch* and *sofa* is not distinguishable from the relation between *sofa* and *couch*.⁴ Thus, we can say that *couch* and *sofa* are synonyms of each other, but cannot say that *cat* and *mammal* are hyponyms of each other. *Cat* is a **hyponym** of *mammal*, and *mammal* is a **hyperonym** of *cat*. Similarly, *meronymy* is a unidirectional term, so that *stanza* is a **meronym** of *poem*, but *poem* is the **holonym** of *stanza*. While Lyons (1977) and others discuss *co-hyponymy* and others write of *co-meronymy*, these two relation types can just as well be considered contrast sets. So, *eyes/nose/mouth* could be considered a contrast set or a set of co-meronyms of *face*, and likewise *sonnet/ballad/ode* are a contrast set or co-hyponyms of *poem*. Subtypes of the lexical relations are defined and discussed in Part II.

Other relations, such as morphological or phonetic relations and undefined relations are indicated by a dash (-). Not all semantic relations are discussed in this book. For example **case relations**, like that between *author* and *book*, are disregarded even though they are relevant to some theories of intralexical organization (e.g., Meaning-Text Theory, see 3.3.3). Some miscellaneous paradigmatic relations are briefly discussed in chapter 6, but the attention here is to those relations that have been central in discussions of lexical semantics.

Some instances of relations, particularly examples of antonymy, seem to have special status, in that their relationships are well known in the culture and seemingly stable. For example, *hot/cold* seems like a better example of antonymy than, say, *steamy/frigid*, even though both pairs indicate opposite extremes on the temperature scale. The *hot/cold* pair and others like it (e.g., *big/little*, *good/bad*, *good/evil*) can be considered **canonical** antonym pairs. These are the kinds of antonyms that automatically follow one another in free word association tasks, or that are collected together in children's books of opposites. The **non-canonical** pairs are less common or more context dependent. The differences between the canonical and non-canonical types are discussed in chapters 2 and 5. For now, note that a complete account of semantic relations must acknowledge both types. The two types are not completely separable – their boundaries are fuzzy and it is not always possible to determine whether a pair is canonical

or not. Certainly, *happy/sad* is canonical, but is *happy/unhappy*? If not (on the grounds that morphologically derived antonyms are of a different category than morphologically unrelated antonyms), then why does *happy/unhappy* seem like a “better” antonym pair than *green/non-green* or *straight/unstraight*? *Wet/dry* is canonical, but is *humid/arid*? *Wet/dry* is certainly a more common pair, but cannot uncommon pairs belong to the canon? Language users can intuitively sort “good” (or prototypical) antonym pairs from not-so-good ones and downright bad ones. A complete theory of semantic relations must account for the continuum of relatedness that is revealed by language users’ judgments of “better” and “worse” examples of these relations (see 2.1.4).

Finally, what do these relations relate? So far, I have referred to them as relations among words, but one might ask if they are not really relations among the words’ denotata. After all, isn’t *hot* the opposite of *cold* because heat and cold are incompatible temperature states? While it is true that referring antonyms describe incompatible referents, there is more to antonymy than this. Defining antonymy as referential incompatibility would mean that *hot*, *boiling*, *steamy*, *warm*, *scorching*, and many other words would be equally appropriate as antonyms for *cold* in any context, since all these words describe states that are incompatible with coldness. We thus need to look in particular at how words, or word meanings, are related – not just at how things in the world are related.

Most lexical semantics texts claim that semantic relations are not really relations among words, but relations among word senses. Some of these texts call these relations *sense relations* (Lyons 1977) or *meaning relations* (Allan 1986) rather than *lexical relations*. I have not adopted these terms for two reasons. Firstly canonical antonyms give cause to wonder whether it is just the senses (and not the words themselves) that are being contrasted, since the contrast can extend to many of the words’ senses. The word pair *hot/cold* can be used to describe a number of sense contrasts: ‘high temperature’/‘low temperature,’ ‘near the searcher’/‘far from the searcher’ (in a hiding game), ‘radioactive’/‘not radioactive’ and so forth. The pair *hot* and *cold* has a connection beyond the relation of their temperature senses, and so we can think of them as lexically related. Secondly, as becomes clearer later in the discussion (especially in 2.1.3 and 2.1.5), senses are not the only determinants of lexical semantic relations. Thus, while some semantic relations are sense relations, *sense relations* describes a narrower range of relations than is discussed here.

A basic question that has not been answered yet is: Where should semantic relations be situated in a model of language competence and language use? Many mental model-makers propose (or assume) that semantic relations among words must be represented in the lexicon, with other knowledge of words (see chapter 3). The other option is to consider semantic relations among words as a form of metalinguistic knowledge. In order to evaluate either of these

possibilities, we must first have a clear idea of what the lexicon is and what it contains, as well as how metalinguistic knowledge is to be modeled. The next two sections concern these issues.

1.3 What is a mental lexicon?

1.3.1 Basic assumptions

If a lexicon is a collection of information about words, then it stands to reason that the mental lexicon is the mental representation of information about words. But what that information is and how it is represented are anything but simple questions. The foremost assumptions about the mental lexicon (henceforth *the lexicon*) in structuralist and generativist theories are (a) that the lexicon is a collection of information that cannot be derived from other information and (b) that this information is self-contained and specifically linguistic. So, lexical information is arbitrary or “idiosyncratic” (Chomsky and Halle 1968: 12) – “a list of basic irregularities” (Bloomfield 1933: 274) for which the grammar cannot account. Furthermore, the lexicon contains information about words, rather than about the things or ideas that words denote. These assumptions lead to two basic problems for determining what information is in the lexicon. First, the identification of irregularities and idiosyncrasies depends upon the theories of grammar and cognition involved. For example, if we assume that part of the lexical representation of a word is a set of semantic features based on semantic primitives (e.g., Katz and Fodor 1963) and that learning the meanings of words involves filling in feature specifications from available referential information (Clark 1973), then we might conclude that semantic relations are derivable from the words’ featural structures and some basic relational rules (Katz 1972). So, for example, *girl* is an antonym of *woman* because antonymy requires similarity of all features but one, and *girl* and *woman* meet this criterion:

<i>girl</i> :	[+ human, – adult, + female]
<i>woman</i> :	[+ human, + adult, + female]

On the other hand, one could theorize that the relations between words are primary, and meanings arise from the networks of relations in the lexicon (Deese 1964, 1965; Quillian 1968; G. Miller 1998a). On this view, *girl* and *woman* are opposed and their relations with other words (*boy*, *man*, *child*, etc.) are represented as arbitrary facts. Each word, then, has a unique position in a relational network, and the meaning of any particular word is derivable from its position in this network. So, if relations are taken to be arbitrary, then meaning is not entirely arbitrary, but if meaning is taken to be arbitrary, then relations are not entirely arbitrary. Thus the question of what is arbitrary cannot be answered in a theory-neutral way. We can only judge which theory is better on the basis

of (a) theory-internal consistency and (b) ability to account for the largest range of data and behavior.

The assumption that the lexicon is part of a modular linguistic faculty brings other problems to the fore, since it can be very difficult (if not impossible) to distinguish between knowledge about words and knowledge about the things the words denote. Modularity is more easily assumed for the grammatical aspects of language. After all, having the ability to make subjects and verbs agree is not obviously necessary for non-linguistic cognition, nor is it clearly analogous to other cognitive abilities. Thus, positing that grammatical knowledge and processes are separate from other mental faculties is not at all controversial in many linguistic camps. Positing that a lexicon with semantic content is separate from other cognitive faculties requires more justification because the boundary between lexical and conceptual information is elusive. We are left with the questions: To what extent does the lexicon include information about words' meanings? Can word meanings be differentiated from concepts?

Putting off the problem of lexical versus conceptual semantics until 1.3.3, I take the modular lexicon assumption as a starting point in evaluating approaches to semantic relations. The angle of argumentation here is to show that a well-defined lexicon cannot be the locus of semantic relation information. Sticking with a strict modular definition, then, the **lexicon** contains all and only information that is: (a) arbitrary (i.e., not derivable from other information) and (b) necessary for linguistic competence. **Linguistic competence**, as it is understood here, is the capacity to produce grammatical and interpretable sentences. Linguistic performance is not mentioned in the definition, since competence in non-linguistic faculties is also relevant to linguistic performance.

The modular lexicon assumption has several benefits. First, by applying a strict definition of what lexical information is (and is not), the content of the lexicon is limited. This gives us the opportunity to disprove the existence of a modular lexicon by systematically showing that each potential piece of lexical information should not be included in the lexicon because either it is not arbitrary or it does not contribute to linguistic competence. The assumption that lexical information is inseparable from conceptual information is not as testable, since the definition of *conceptual information* is less constrained than that of *lexical information* in modular models of the mind.

Ease (and reliability) of argumentation is not the only reason to assume that lexical information is separable from conceptual information. An alternative is to consider lexical information as a subtype of conceptual information. So, in addition to representing facts about a concept TABLE like 'this thing is furniture; it has a flat top surface; I can put things on it,' part of the concept TABLE would be 'the word for this concept is *table*; it's pronounced [tebɪ]; it's a noun.' But there are reasons to believe that lexical forms are not represented as parts of (non-lexical) concepts (see also Clark and Clark 1977). For one,

lexical information must be accessible to different types of mental processes than conceptual information is. Since lexical material must all be accessible to syntactic and phonological rules and non-lexical material need not, it is reasonable and expedient to treat lexical matter differently than other conceptual information. The fact that we can fail to make the association between things that we recognize and words that we know for those things also indicates that our means of storing and/or accessing the name of a thing is not the same as our means of storing and/or accessing other knowledge about the thing. While striking examples of this are seen in aphasia and acquired anomia, we all experience this problem sometimes in tip-of-the-tongue syndrome. In this case, you have complete access to the concept, since you can picture it, reason about it, and describe it (“You know, it’s like a puppet, but it has strings”), but you are not able to access its name. Other evidence for the separation of lexical and conceptual information is related to the lack of one-to-one relationships between words and concepts. If names for things were part of our conceptual representations of those things, then we should have unambiguous mapping between name and thing. Words can be used to indicate more than a single concept, however, and the name that we attach to a thing may vary by context. In the first case, the word *knife* can refer to things like scalpels, daggers, butter knives and letter-openers (Cruse 2000a); in the second, a single kind of furniture may be referred to by a variety of terms like *table*, *bedstand*, and *chest of drawers*. We need means to mix and match names and concepts, rather than static concept-word association.

Thus, knowledge of words is a different type of knowledge than knowledge of things. (But words can also be things. See 1.4.) These two types of knowledge interact in the processes of language production and comprehension. The contents and structure of the lexicon are described in the next subsections.

1.3.2 *Lexical items and lexical entries*

While *word* is the word that has been used thus far to indicate lexical items, it is certainly not an accurate one. The lexicon must contain both linguistic expressions that are greater than words and ones that are smaller than words because they too may be non-predictable in their mapping of form to meaning. Non-compositional phrasal expressions, such as *throw up* or *paint the town red* and arguably morphemes, such as *-ness* or *pre-*, are also to be included in our definition of **lexical item**, or **lexeme**.

Some linguists debate whether words are represented in the lexicon as self-contained units. Sinclair (1998) and Weigand (1998a) argue that idiosyncratic form-meaning associations involve larger expressions than words, and thus the lexicon is not a simple catalogue of words (and other non-compositional

expressions) that can enter into productive grammatical processes. For example, Weigand (1998a) notes that different adjectives may indicate the same property when collocated with different nouns, and it is not particularly predictable which nouns go with which adjectives. So, for example, the meaning that she expresses as SHAPE/DENSE is communicated by *heavy* in *heavy traffic* and by *thick* in *thick forest*. She concludes that the lexicon must contain complex lexical items that allow for the arbitrary nature of these collocations and the particular meanings they involve. This is a relevant point in accounting for the syntagmatic semantic relations between adjective and noun, but it is less clear that these multi-word constructions are relevant to the study of paradigmatic lexical relations, which stereotypically involve the relation of word-length items. That words are interesting in and of themselves becomes clear when we look for antonyms of these phrases. The opposite of *heavy traffic* is *light traffic*. The opposition between these phrases relies upon the canonical opposition of the adjectives in those phrases, as seen in various uses of *heavy* and *light*, whether they are used to indicate density (*heavy/light rain*), weight (*heavy/light luggage*), richness (*heavy/light meal*), oppressiveness (*a heavy/light mood*), and so forth. Since paradigmatic relations are of primary concern here, words continue to be important to the notion of *lexical item*. So, while not precluding the existence of multi-word lexical items, words should be considered as independent units in those phrases for the purposes of some paradigmatic relations (not to mention morphological processes).

A lexical item in the lexicon is an abstract representation that is instantiated as a **lexical unit** in language use (Cruse 1986), which has a particular form and a particular sense. So, for example, *highest* in the phrase *the highest note in the song* and *high* in *I threw the ball high* are both lexical units instantiating the lexical item *high*.

Unlike inflectional variations (e.g., *high* → *highest*), morphological derivations (e.g., *high* → *highness*) must often be represented as lexical items, either because preference for a particular form is arbitrary (e.g., *typist* over *typer*) or because the meaning is not compositional (e.g., *highness* as referring to a royal person). Whether non-compositional expressions are included in the lexicon is a matter of some debate. Bybee (1985, 1998) has argued that some morphologically derivable words are recorded in the lexicon. A semantically compositional word might be included in the lexicon because its frequency makes it more conveniently stored in the lexicon than derived anew again and again, or it might be included because it was through this form that we learned the base form of the word. In a similar vein, Jackendoff (1997) has argued that conventional yet semantically compositional strings of words (such as nursery rhymes, song lyrics, clichés) should be treated as lexical units. Again, this makes sense because such strings are obviously not composed anew each time

they are uttered. While such arguments may indicate that the lexicon contains non-arbitrary strings of words, they are in fact arbitrary in that they have been conventionalized as a particular form.

Lexical entry describes the collection of information about a lexeme that is included in the lexicon. At the very least, a lexical entry must include phonemic information about the item and some means for mapping the phonemic form to a meaning. The entry might also include grammatical category and other unpredictable information.⁵ Some treatments of the lexicon (especially processing-oriented models, e.g., Levelt 1989) distinguish the representations of a word's form (i.e., phonemic information) from its grammatical and semantic content. The latter is termed the **lemma**.

This abundance of terms (*word*, *lexeme*, *lexical item*, *lexical unit*, *lemma*) brings us back to the question of what semantic relations relate. While I have been calling them "relations among words," would they more accurately be described as relations among lemmata or relations among lexical items or units? *Word* is an insufficient description of the items that can be related. Like words, bound morphemes and some idiomatic phrases sometimes have antonyms or synonyms. For example, we might say that *in-* and *un-* are synonymous in contexts like *inedible=unedible*, and that the phrase *black and white* is the antonym of the word *color* when discussing photography or film. So, semantic relations may relate lexically represented expressions other than words. On the other hand, *lexical item*, *lemma*, and *lexical unit* are also insufficient, for reasons that are discussed in chapter 2. I therefore continue to refer to the items in semantic relations as *words*, with the understanding that claims made here about words may be extended to non-word lexical items.

1.3.3 *The dictionary metaphor and the nature of meaning*

The term *mental lexicon* is an artifact of a pervasive metaphor for the mental representation of vocabulary: that of a dictionary in our heads. Of course, the lexicon-as-dictionary metaphor fails on many counts, since the mental lexicon and the printed dictionary represent some different types of information about words. The mental lexicon includes phonemic structures, but print dictionaries record quasi-phonetic forms. My mental lexicon must record that *glad* is not to be used in prenominal position (since **the glad person/occasion* is ungrammatical), but the particular dictionary beside me (*American Heritage*, 4th ed.) does not. Instead, it gives the word's etymology, which is irrelevant to the mental lexicon. While dictionaries list meanings of words, they cannot list all the meanings for which we use those words, since the potential semantic uses of a word are without limit (Nunberg 1978). Unlike dictionaries, the mental lexicon cannot afford an arbitrary separation of definitional and "encyclopedic"

meaning, nor are its sense divisions the same as those in dictionaries. These last two problems are discussed in turn below.

Hand-in-hand with the lexicon-as-dictionary metaphor is the conceptual-knowledge-as-encyclopedia metaphor. On this view, the lexicon should include only definitional (core semantic) information, leaving encyclopedic (conceptual) information in the realm of what we know about things, rather than what we know about words for things. So, for example, that dogs are used as sled-pullers might be something I know about dogs, but is not part of the meaning of *dog*. One means for differentiating definitional and encyclopedic meaning would be to assume that word meanings can be defined on the basis of necessary and sufficient conditions (Katz and Fodor 1963). However, most of our everyday content words cannot be defined by necessary and sufficient conditions, as Wittgenstein (1958) showed for *game*. Further experimental work (e.g., Labov 1973; Rosch 1973, 1975, 1978) has shown that word meaning seems to be organized around conceptual prototypes. Prototype approaches to meaning blur (or erase) the line between the definitional and the encyclopedic and call into question the separation of linguistic and conceptual semantic knowledge (see Taylor 1995). If we accept the prototype approach to meaning, this blurring between lexicon and encyclopedia is the first problem for the lexicon-as-dictionary metaphor.

Another means for differentiating definitional and encyclopedic meaning is to attempt a separation of linguistic and non-linguistic information. In such an approach, aspects of meaning that have grammatical reflexes and thus are language specific (rather than generally conceptual) must be represented at some linguistic (lexical) semantic level (Gruber 1983; Pinker 1989). Representing these aspects of meaning would not involve representing full senses, and so conceptual meaning would still be crucial in comprehending utterances. Some of this lexically represented semantic information is potentially relevant to semantic relations. For example, the thematic relations for *buy* and *sell* are converse, in that the role assigned to the subject position for *buy* is assigned to the object position for *sell* and vice versa. This converseness of thematic structure could be considered to be related to the status of *buy/sell* as converse antonyms. But if the semantic information in the lexicon is only that which interacts with the grammar, then the lexicon would not encode enough semantic information to account for all semantic relations. For example, *give* and *sell* might have the same thematic structure – but that does not make them synonyms. Similarly, the grammatical information alone would give us no way to distinguish between *give* and *sell* as potential antonyms for *buy*. So, the grammatically relevant information in lexical entries is not sufficient for determining semantic relations. Sense information is needed as well – but it is far from clear that senses are included in lexical entries, as discussed below.

The second problem for the lexicon-as-dictionary metaphor is the fact that lexical items can map to many different concepts, and thus be **polysemous**, but there is no principled limit to a word's polysemy. A dictionary deals with polysemy by listing a small number of senses for any word in an entry, and some lexical semantic theories have treated polysemous words as having multiple senses within lexical entries (e.g., Ullmann 1957; Katz and Fodor 1963). So, like a dictionary, the lexical entry for *horseradish* might list three possible meanings: a type of plant, the vegetable that consists of the root of that plant, and a condiment made from that vegetable. But listing meanings in the lexicon is doomed to failure since, as Nunberg (1978) has argued, the number of usable senses for any lexical item is limitless. Nunberg's argument is based on the observation that different criteria for establishing reference can be invented and used within the particulars of a context – so nonce meanings are possible and not uncommon. Nunberg uses the example of *jazz*, but *tea* illustrates the point as well. *Tea* can refer to a drinkable herbal infusion or the prepared herbs for making such an infusion. It can also refer to a cup- or glass-sized portion of that infusion, as in *I'd like a tea, please*. It can also refer particularly to a hot version of this drink (in contrast to *iced tea*), but in the southern United States it refers to the iced version, in contrast to *hot tea*. It is also sometimes used to refer to certain types of tea (especially those with caffeine), so that we may contrast it to others, as in *I can't drink tea after supper – just herbal tea*. Such conventional uses are probably countable in number, and some may be excluded from any particular English speaker's lexicon because they are not part of that person's dialect. But even within a single language user, the range of concepts that a lexical item indicates is not necessarily limited or static. For example, let us say that in South Africa I grew to like rooibos tea and that I visit Nancy in New York who asks *Would you like some tea?* Now, knowing that Nancy has probably never heard of rooibos, I assume that when she says *tea*, rooibos is not a member of the set of things that she intends to refer to, so I reply, *No, I don't care for tea*. For the purpose of this exchange, the sense I use for *tea* does not include rooibos, but in another context I may refer to rooibos as *tea*, as in *The only tea I like is rooibos*. I leave it to the reader to imagine other contexts in which a speaker might use a sense of *tea* that denotes all teas but chamomile or only peppermint tea. The point is that the category that the speaker intends to refer to with the word *tea* (and that the audience may identify when the speaker uses *tea*) shifts with the speaker's knowledge and expectations of the context. Thus, the number of possible senses of *tea* that may be reasonably intended and understood is limited only by the number of possible combinations of beliefs that the speaker and hearer have about the world and the situation in which the utterance is made. Thus, one cannot take an inventory of a word's senses. Instead, a word's sense in any particular context is the result of some implicit negotiation between the members of

the talk exchange along with beliefs about how that word is conventionally used.

Since words have no fixed number of senses, it is untenable to claim that lexical entries explicitly represent all of a word's senses. One solution is to assume that the various meanings attributed to any particular word are, at some level, illusory. Weinreich (1963), for example, claims that the many uses of the verb *take* indicate not that the word has many meanings, but that it is semantically nearly empty. While Weinreich contrasts such cases to other cases of true polysemy, Ruhl (1989) applies semantic underspecification generally, arguing that all words are monosemous. Another possible approach is to assume that polysemous words have a single sense that is part of the lexical entry and that other senses are derived by lexical rules (e.g., Jackendoff 1976; Pustejovsky 1995; Copestake and Briscoe 1995).⁶ In this case, the lexically represented sense information may or may not be structurally simpler than the other possible senses, but it is a starting point from which semantic information may be added or deleted in order to build new senses. Nunberg holds that knowledge of word meanings is a type of knowledge of "the collective beliefs of the speech community" (1978: iii). In this way, knowledge of word meaning is not simply linguistic knowledge – it interacts with and is constrained by knowledge of how members of the speech community communicate and the beliefs they hold about the words and the objects and situations they designate. Recognizing or intending a particular sense for a word is possible because we are aware of some conventions for the word's use, and we have pragmatic means for creating recognizable new senses for a word.⁷ So, we have three possibilities: (a) multiple senses are illusory (each word has only one sense), (b) additional senses are derived from a basic sense representation, or (c) no senses are basic, but instead meanings are generated through pragmatic knowledge.

If we follow Lyons' (1977) claim that semantic relations relate senses of words and we understand *sense* as Nunberg's context-dependent conditions on reference, then it is impossible to represent semantic relations in the lexicon. In order to represent sense relations in the lexicon, all the senses would have to be represented in the lexicon. In the monosemy approach, all the senses are listed in the lexicon, since every word only has one sense. But words can have more than one hyperonym (or antonym or synonym), and multiple hyperonyms need not be synonymous with each other. For example, *measurement* and *spoon* can both be hyperonyms for *teaspoon*. In order to account for the fact that the same word is involved in incompatible semantic relations, the monosemic approach either has to declare all such examples to be cases of **homonymy** (i.e., each meaning is associated with a different lexical entry, as would be the case for unrelated words like *bat* 'flying mammal' and *bat* 'wooden club for games'), or it has to claim that semantic relations are not sense relations, but relations that arise through the conditions in which a word is used. The

pragmatic approach (Nunberg) would expect that semantic relations are either part of the community's linguistic convention (i.e., not sense-related, but an arbitrary fact about language) or they would be derived from the contextual conditions of a word's use. The lexical-rule approach, in which new senses are generated, would allow for some semantic relations to be represented in the lexicon (but the nature of sense representations would make such information redundant, see 3.2). Since other senses are derived, relations among them would also have to be derived.

Where I need to make assumptions in the next chapters about the nature of word meaning, the following are assumed. Firstly, words are polysemous: they can be associated with more than one sense (i.e., I reject the monosemy solution). Secondly, a **sense** is the set of conditions on a word's denotation. Connotation is a separate matter. So, for example, if I use *green* in *the green grapefruit* or *the green army fatigues*, the two uses do not have different senses even though (a) in the first (but not the second) case it signals unripeness (and hence connotes sourness) and (b) the referent (i.e., the shade of green) is different in the two cases. In both cases, the same conditions of reference applied, namely, whether the color of the item fits within a certain range of hues. Thirdly, while some semantic information may be represented in the lexicon, senses are not represented intralexically. A sense in toto is composed from whatever semantic information is specified in a lexical entry, the information (about the denotation of the word) that the word maps to in the conceptual realm, and contextual information. Sticking with the example *green*, then, the lexical entry might need to represent some semantic information, like that it describes a gradable property (and hence can be used in comparatives).⁸ Through the context, we figure out how to map the word to a concept or set of concepts, such as the concept of a color (as in *the green grass*), of a thing that has that color (as in *I've run out of paint, give me some more green*), or of some specific quality associated with that color (as in *the green intern, green politics*). The context also allows us to determine the boundaries of the denotation, for instance whether for our purposes *green* excludes OLIVE GREEN or refers only to FOCAL GREEN. So, in essence, I assume senses to be dynamic, and assume that the fixed mental representations of semantic information (lexical or conceptual) allow for adaptation to the requirements of a particular context. Senses that seem basic to words are usually those that require the fewest contextual cues or lexical/conceptual processes and/or that refer to more prototypical exemplars of the concepts involved. So, 'green-colored' is a more basic sense for *green* than 'green paint' because the latter includes the meaning of the former and was presumably derived through a process that lets color words stand for things that are the color that the word describes and then applied to a particular context (in which it specifically means 'green paint' rather than 'thing that is green'). A sense of *green* in which FOCAL GREEN is the

defining prototype is more basic than a sense that only refers to shades of chartreuse.

1.3.4 *The thesaurus metaphor*

Besides the dictionary metaphor, some models of the lexicon instead (or also) employ what I will call the **thesaurus metaphor**. In such relational models of the lexicon, lexical entries (or senses within them) are cross-referenced with or linked to other lexical entries (or senses within them). Thus, the information that *high* is the antonym of *low* is represented in the lexicon in such models, and may serve as the means by which they represent the type of basic semantic information about the words that is arbitrary and stable. In such models, since *high* and *low* are in an antonymic relation, they are constrained to always denote the opposite ends of whatever scale they describe. The specifics of such models, and their advantages and disadvantages, are discussed in chapter 3. For the time being, it is worth noting that dictionary and thesaurus models exist on a continuum, and that the two ends of the continuum represent **componentialism** and **holism**. On the extreme componentialist end are dictionary models that are not relational at all, in which the lexical entries are unordered with respect to each other. The assumption of an unordered lexicon is often found in generative linguistics (e.g., Chomsky 1965; di Sciullo and Williams 1987). At the holistic extreme are thesaurus models that contain no definitional information in the lexical entries, but instead expect meaning to arise through the relations of words in the language's lexical network. This position is best represented in some computational models (e.g., Quillian 1968 and early versions of WordNet; see 3.1.5 and 3.4.2, respectively), but it can also be seen in some European structuralist linguistic positions (see 3.1.2) and in philosophical approaches to meaning (see 3.1.1 and 3.4.1). Many models of the lexicon, including several structural semantic models, fall somewhere between these extremes, showing some preference for definition or relation as the basic means for representing semantic information, but not excluding the other means (see 3.3). I use the term **associationist** to refer to any approach (holist or mixed) that embraces the thesaurus metaphor and treats semantic relations as fixed mental representations.

1.4 Words, concepts, and concepts of words

The linguistic faculty involves two types of knowledge: lexical and grammatical. Each of these types of knowledge is relevant to our use of words (and to our linguistic competence in general), but in different ways. The lexicon represents facts about particular linguistic expressions (i.e., the lexical items), such as the facts that there is a word *night* that is pronounced [najt] and that word is a noun

(see n. 5). The grammar consists of rules for making new linguistic expressions out of the raw materials that the lexicon supplies.

Like linguistic knowledge, conceptual knowledge involves representations of arbitrary concepts and rules for using those concepts, allowing for the generation of new concepts. Arbitrary facts are things that one has to experience in order to learn. So, for example, I know that mature kiwifruits are kind of hairy because at some point I was introduced to these facts and made them part of my KIWI concept. While there may be breeds of hairless kiwis or explanations for the hairiness of kiwi skins, I do not know about them. The information that kiwis are hairy is, to me, an arbitrary fact about kiwis.

The conceptual system is accessible to principled processes for making inferences about those arbitrary facts and for constructing new concepts (as for ad hoc categories, Barsalou 1983). For instance, in my long-term memory I have concepts of all the clothes I now own. In imagining A GOOD OUTFIT TO WEAR TOMORROW, I use those conceptual representations of my clothes along with my conceptual representations of OUTFIT (color and style coordination, coverage of the body from at least chest to thigh), tomorrow's weather forecast, my plans for tomorrow (and what clothing is appropriate to them), and so forth. I thus build a concept of a new category (A GOOD OUTFIT FOR TOMORROW) using extant concepts and principles for developing new concepts based on old ones.

As in lexical representation, the mechanics of conceptual representation are subconscious. When I imagine a kiwi, I am aware that I have a conceptual representation of KIWI, but I do not see my conceptual representation of KIWI. My imagined kiwi is the result of applying some processes (recall, mental imaging, and whatever sub-processes they involve) to that particular concept. Presumably, my concept for KIWI is schematized so that it can engage in a number of different cognitive processes, such as categorization, recognition, recall, and mapping to a lexical item.

If the conceptual system represents knowledge of the world, then it stands to reason that it includes representations of language, since language is part of the world that we know. So, we need to differentiate linguistic knowledge (the grammar and lexicon) from the metalinguistic knowledge represented in the conceptual system. Thinking about words is a metalinguistic endeavor, since we do not have direct, conscious access to the structures in the lexicon. Thus, the objects we reflect upon when reading a dictionary or writing a linguistics book are our perceptions of words, and the objects with which we are doing this reflecting are the concepts we have of the words. So when we think about words *in the lexicon*, we are always at least one step removed from the object of our study. We make inferences about words, rather than observations of lexical entries.

Like other conceptual knowledge, our knowledge about words may be derived or remembered. So, for example, I have gathered from experience (and now

remember) that certain words offend my mother and others offend my father, that *anil* is a common word in American crossword puzzles, and that my brother used to metathesize the first two sounds in *spaghetti*. None of these facts is relevant to my competent grammatical and meaningful use of these words (although they may affect where I choose to use them), so there is no reason to believe it is part of my lexical representation of the words.

Other conceptual knowledge of words can be derived extemporaneously. To use a phonological example, when thinking about the word *boot*, I know it has three sounds, but there is no reason to believe that this fact is intralexically or conceptually represented. What is more likely represented in the lexicon is the fact that the word has the sounds [b], [u], and [t], in that order. In observing the word, I have made the inference that its sounds are three in number. This fact is not relevant to my actual use of *boot* as a piece of language – there are no phonological rules that apply only to words with three phonemes. Thus, in reflecting upon the word *boot*, I perceive the word and conceptualize it (or make use of my conceptualization of it).

Some word-concepts may be completely ad hoc (see Barsalou 1983). If I have used a word for years without thinking about it and without any incident that might make it remarkable to me, it might exist as a lexical entry in my mind, but not as a concept. If someone then asks me “what do you know about *only*?” I could reflect on it – thereby conceptualizing it – and pull together some facts about *only*: it rhymes with *lonely*, it has four letters, it can be an adjective (as in *an only child*) or an adverb, sometimes it can be interchanged with *just*, it is not easy to say it backwards, and so forth. I come to these conclusions and thus build my concept for *ONLY* by mentally rehearsing its use and noticing things about its pronunciation, spelling, meaning, and grammatical properties.

This differentiation of lexical and conceptual representation of words is probably not controversial, but acknowledgment of the distinction between the two types of word-knowledge is necessary in accounting for some phenomena that traipse the lexical-conceptual divide. Conflicts between metalinguistic beliefs and facts about language use indicate that our metalinguistic knowledge of words is extensive, if not always perfect. For example, corpus studies often reveal that the ways in which we use words are quite different from our beliefs about how the words are used. For example, if you ask someone (preferably not a corpus linguist) for the most common meaning of *run*, they will probably tell you that it has to do with moving swiftly by foot. But checking the uses of *run* in a corpus, very few of the tokens have a locomotive meaning – most mean something more like ‘manage,’ as in *run a business* (Michael Barlow, 1999 LSA Institute lecture). The mismatch between fact and belief indicates that we have (or create, in order to answer the question) some conceptual representation of *run*, and the prototype for that concept is the version of *run* that indicates a certain form of animal (likely human) movement.

The argument in chapter 2 makes use of the distinction between lexical and conceptual knowledge of words, showing that this distinction is necessary in accounting for semantic relations among words.

1.5 Summary and next steps

In this chapter I have defined paradigmatic semantic relations among words as the topic of study. The goal in approaching this topic is to create a psychologically viable model of how such semantic relations are acquired and used. Thus, the overarching question is: How are semantic relations mentally represented? The perspective taken here is pragmatic, in that it is assumed that we must examine the use of words in context in order to appreciate how they might be represented in the mind.

There are two basic types of answer to the question of how semantic relations are mentally represented. Either relations among words are represented directly (thesaurus-style), as facts that we know, or they are derived by some sort of relational rules that take into account other facts that we know (such as the words' meanings). In order to test the idea that semantic relations might be directly represented in the lexicon, I have defined the lexicon precisely enough to make clear the implications of this idea. Thus, the lexicon is defined as the repository of arbitrary facts about language – that is, linguistic information that is not derivable from other knowledge. The problems of polysemy and of distinguishing the lexical-encyclopedic divide necessitate that the lexicon be semantically impoverished. Finally, two types of mental representations of words were noted: (linguistic) lexical entries (represented by italics: *lexical item*) and (metalinguistic) conceptual representations of words (in italic small capitals: *LEXICAL CONCEPT*).

The remainder of Part I concerns semantic relations in general. Chapter 2 describes the properties of semantic relations for which a theory of semantic relations must account. It then introduces a pragmatic, **metalexical** account for paradigmatic semantic relations based on the assumptions outlined in this chapter. Chapter 3 provides a multi-disciplinary survey of other approaches to semantic relations and how they compare to the metalexical approach. Part II (chapters 4–6) concerns particular relations in detail, with chapters on synonymy, antonymy and contrast, and hyponymy, meronymy and other relations. In the final chapter, I summarize the foregoing, reconsider some problems for the metalexical account, review compatible models of the lexicon, and suggest further questions for investigation.