

Cambridge University Press

0521780675 - Semantic Relations and the Lexicon: Antonymy, Synonymy, and Other
Paradigms

M. Lynne Murphy

Excerpt

[More information](#)

Part I

Paradigmatic relations, generally

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

Cambridge University Press

0521780675 - Semantic Relations and the Lexicon: Antonymy, Synonymy, and Other Paradigms

M. Lynne Murphy

Excerpt

[More information](#)

4 Paradigmatic relations, generally

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

Cambridge University Press

0521780675 - Semantic Relations and the Lexicon: Antonymy, Synonymy, and Other Paradigms

M. Lynne Murphy

Excerpt

[More information](#)

Why lexical relations?

5

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 fulfills 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

Cambridge University Press

0521780675 - Semantic Relations and the Lexicon: Antonymy, Synonymy, and Other Paradigms

M. Lynne Murphy

Excerpt

[More information](#)

6 Paradigmatic relations, generally

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

Cambridge University Press

0521780675 - Semantic Relations and the Lexicon: Antonymy, Synonymy, and Other Paradigms

M. Lynne Murphy

Excerpt

[More information](#)

Why lexical relations?

7

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

Cambridge University Press

0521780675 - Semantic Relations and the Lexicon: Antonymy, Synonymy, and Other Paradigms

M. Lynne Murphy

Excerpt

[More information](#)

Why lexical relations?

9

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 (<)

Cambridge University Press

0521780675 - Semantic Relations and the Lexicon: Antonymy, Synonymy, and Other Paradigms

M. Lynne Murphy

Excerpt

[More information](#)

10 Paradigmatic relations, generally

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

Cambridge University Press

0521780675 - Semantic Relations and the Lexicon: Antonymy, Synonymy, and Other Paradigms

M. Lynne Murphy

Excerpt

[More information](#)

Why lexical relations?

11

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