CHAPTER ONE

What Are Word Meanings?

Learning Objectives

- Understand that word meanings function as publicly shared concepts for social coordination.
- Understand that word meanings consist of features that are studied in different ways by linguists and psychologists.
- Describe how frames organize the features that make up word meanings.
- Describe how prototypes provide cognitive reference points for word meanings.
- Describe how mental models enrich word meanings.
- Describe how background situations are associated with word meanings.
- Describe how traditional semantic relations (synonymy, antonymy, and taxonomy), thematic relations, and collocational relations influence word meanings.
- Describe how linguistic contexts constrain the processing of word meanings.
- Understand that there are both similarities and differences between languages regarding word meanings.

1.1 Introduction

Most word meanings are complex packages of information. This chapter surveys many aspects of them, but it is restricted to those aspects that have received considerable attention in neuroscience, as described in Chapters 3–7. Among the aspects that have not yet been carefully studied in neuroscience, two are worth mentioning briefly here at the outset.

First, words vary with respect to their internal architecture (Lieber, 2021). Some words, like *book*, *walk*, and *think*, exist as wholes that can't be broken down into smaller meaningful parts. In contrast, other words, like *reconnected*, *unpacking*, and *disagreement*, can be divided up into chunks that each have their own sense or function. For example, *disagreement* has three elements: (1) *agree*, which is the stem; (2) *dis-*, which is a negative prefix that also appears in words like *dislike*, *disobey*, and *disprove*; and (3) *-ment*, which is a nominalizing suffix that also appears in words like *achievement*, *excitement*,

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

and *improvement*. So far, not much is known about how our brains derive the composite meanings of internally complex words from the separate meanings of their parts.

Second, although many semantic properties of words are effectively invisible to grammar, others are quite relevant (Frawley, 1992). For example, adjectives fall into multiple classes according to general semantic properties such as size (*big*, *little*, *large*, *small*), dimension (*tall*, *short*, *thick*, *thin*), and color (*blue*, *green*, *yellow*, *red*). When two or more adjectives are placed before a noun, what typically determines their linear order are these general semantic properties, not the specific ones that distinguish between the adjectives in each class. Thus, size adjectives tend to precede dimension adjectives, which is why it's better to say *I saw a big tall building* than *I saw a tall big building*, and dimension adjectives tend to precede color adjectives, which is why it's better to say *I have a thick blue towel* than *I have a blue thick towel*. Similar distinctions between grammatically relevant and irrelevant semantic properties have also been identified for nouns and verbs. As yet, however, very little has been learned about how such distinctions are implemented in the brain.

On the other hand, as noted above, substantial progress has been made in elucidating the cortical underpinnings of many other aspects of word meanings, and the purpose of this chapter is to provide an initial overview of these topics, highlighting insights from linguistics and psychology. Before delving into this material, though, the first section below sets the stage by addressing the social functions of word meanings.

1.2 Word Meanings Are Publicly Shared Concepts for Social Coordination

Word meanings are part of an intricately organized conceptual system that is essential for a wide range of mental activities that go far beyond language. Most importantly, concepts, whether lexically encoded or not, are coherent, relatively stable units of long-term knowledge that allow us to classify certain phenomena as being unique tokens of more general types – for example, a member of the category of dogs. As Hofstadter and Sander (2013, p. 15) poetically put it, "nonstop categorization is every bit as indispensable to our survival in the world as is the nonstop beating of our hearts. Without the ceaseless pulsating heartbeat of our 'categorization engine,' we would understand nothing around us, could not reason in any form whatever, could not communicate with anyone else, and would have no basis on which to take any action."

Word meanings are special, however, because they embody concepts that, by virtue of having conventional labels, are shared by most people in the language

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

community. The primary utility of such publicly shared concepts is that they allow people to manage problems of social coordination (Enfield, 2015, 2022). Whatever goal two conversation partners may have – whether it be to negotiate, commiserate, fight, or flirt – their words create mutually understood and jointly attended versions of reality. But because word choices are always subjective, these lexically invoked versions of reality inevitably reflect the speaker's own perspective. As Tomasello (1999, pp. 8–9) points out, "in different communicative situations one and the same object may be construed as a dog, an animal, a pet, or a pest; one and the same event may be construed as running, moving, fleeing, or surviving; one and the same place may be construed as the coast, the shore, the beach, or the sand . . . ."

The inherently perspectival nature of word choice implies that the way a speaker portrays a scene has psychological consequences for the hearer. In a famous experiment, Loftus and Palmer (1974) showed different groups of people the same video of two cars colliding and later tested their memory by asking them slightly different questions: "How fast were the cars going when they bumped into each other?" versus "How fast were the cars going when they smashed into each other?" Compared to the people in the "bumped" condition, those in the "smashed" condition not only estimated a higher speed, but were also more likely to say, incorrectly, that there was broken glass on the ground. The different labels distorted people's memories of the same event in different ways by encouraging different construals of it. More generally, much of everyday discourse involves speakers selecting certain words rather than others in order to influence how hearers think about the world. Such rhetorical strategies usually have a helpful, cooperative intent. However, Enfield (2022, p. 198) observes that in adversarial situations "our words are bids to set the terms that suit us best, not necessarily the terms that get us closest to the truth, [as shown by] ideologically charged disputes over *riot* vs. *protest* or *terrorist* vs. *freedom fighter*."

1.3 Features

The most common way to characterize word meanings is to decompose them into their component parts, which are often called features. How does this work? Traditionally, linguists have relied mainly on intuitive structural analysis. To take a few classic examples, it's reasonable to assume that *kill* means <cause to die> and *bachelor* means <an unmarried man>. Similarly, it seems obvious that the meaning of *bird* includes features like <feathers>, <wings>, <beak>, <flies>, and <lays eggs>. This intuitive approach not only illuminates the content of individual word meanings but also reveals similarities between them. For instance, *hop*, *skip*, *walk*, and *run* all share the feature <move>; *eat*,

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

dine, *devour*, and *drink* all share the feature <ingest>; and *mother*, *sister*, *waitress*, and *mare* all share the feature <female>.

To achieve greater reliability, however, psychologists usually prefer to decompose word meanings in more experimentally controlled ways that allow them to infer statistical regularities from data provided by groups of participants. In one approach, researchers first ask participants to list lots of features of certain words, and then they calculate how frequently those features are generated (Buchanan et al., 2019; Devereux et al., 2014; McRae et al., 2005; Vinson & Vigliocco, 2008). For example, Table 1.1 shows the features listed for *moose* and *knife* together with their frequencies. Studies of this kind have led to the creation of normative databases that allow researchers to conduct further investigations, such as exploring how words with shared features form categories (e.g., animals, vegetables, furniture) and how specific features affect people's performance on tasks like word recognition and picture naming.

Table 1.1 Semantic features generated by 30 participants for the words *moose* and *knife*, together with the production frequency and classification of each feature. (From McRae et al., 2005, pp. 555–556.)

Concept Name	Feature	Production Frequency	Classification
moose	is large	27	visual – form and surface
	has antlers	23	visual – form and surface
	has legs	14	visual – form and surface
	has four legs	12	visual – form and surface
	has fur	7	visual – form and surface
	has hair	5	visual – form and surface
	has hooves	5	visual – form and surface
	is brown	10	visual – color
	hunted by people	17	function
	eaten as meat	5	function
	lives in woods	14	encyclopedic
	lives in wilderness	8	encyclopedic
	an animal	17	taxonomic
	a mammal	9	taxonomic
	an herbivore	8	taxonomic
knife	has a handle	14	visual – form and surface
	has a blade	11	visual – form and surface
	made of steel	8	visual – form and surface

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

Table 1.1 (cont.)

Concept Name	Feature	Production Frequency	Classification
	made of metal	7	visual – form and surface
	made of stainless steel	5	visual – form and surface
	is shiny	5	visual – form and surface
	used for cutting	25	function
	used for killing	7	function
	used by butchers	5	function
	is sharp	29	tactile
	is serrated	8	tactile
	is dangerous	14	encyclopedic
	found in kitchens	8	encyclopedic
	used with forks	6	encyclopedic
	a weapon	11	taxonomic
	a utensil	9	taxonomic
	a type of cutlery	5	taxonomic

In another approach, researchers first ask participants to provide scalar ratings of the degree to which certain words encode certain types of features, and then analyze the results to draw generalizations about the meanings of individual words as well as selected sets of words (for a survey of normative databases derived from this approach, see Gao et al., 2023). During the past few decades, many types of features have been examined – for example, concreteness (i.e., the degree to which a concept can be experienced through the senses), socialness (i.e., the degree to which a concept has social relevance), and arousal (i.e., the degree to which a concept evokes emotion). Some studies that focus on just one type of feature have acquired ratings for huge samples of words. For instance, Brysbaert et al. (2014) collected concreteness ratings for 40,000 words. Other studies have included multiple types of features, but for practical reasons most of them have been restricted to smaller samples of words. For instance, Binder et al. (2016) posited 65 kinds of "experiential attributes" of concepts (see Table 1.2) but only obtained ratings of these attributes for 535 words (see the three examples in Figure 1.1). Despite the relatively small size of this sample, however, the words covered many conceptual domains, and the multidimensional ratings for them constitute a valuable resource, as we will see in later chapters. Finally, it's notable that Lynott et al. (2020) recently reported

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

Table 1.2 Binder et al.'s (2016) componential approach to conceptual analysis includes 65 experiential attributes from 14 categories. See Figure 1.1 for application to the meanings of the words *egg*, *bicycle*, and *agreement.* (Adapted from Binder et al., 2016, p. 133.)

Category	Attributes
Vision	vision, bright, dark, color, pattern, large, small, motion, biomotion, fast, slow, shape
Somatic	touch, temperature, texture, weight, pain
Audition	audition, loud, low, high, sound, music, speech
Gustation	taste
Olfaction	smell
Motor	head, upper limb, lower limb, practice
Spatial	landmark, path, scene, near, toward, away, number
Temporal	time, duration, long, short
Causal	caused, consequential
Social	social, human, communication, self
Cognition	cognition
Emotion	benefit, harm, pleasant, unpleasant, happy, sad, angry, disgusted, fearful, surprised
Drive	drive, needs
Attention	attention, arousal

a particularly impressive megastudy. They managed to gather ratings of 11 types of sensory/motor features – six perceptual modalities (visual, auditory, haptic, gustatory, olfactory, and interoceptive) and five action effectors (head, mouth, torso, hand/arm, and foot/leg) – for a total of 39,707 words, which approximates the vocabulary size of an average native English speaker in their mid-30s (Brysbaert et al., 2016). The diversity of results is illustrated in Figure 1.2.

All these methods of identifying the building blocks of word meaning have yielded copious insights, but they also face several challenges. Here I'll mention just three. First, seemingly straightforward definitions often turn out to be inadequate. For example, Geeraerts (2010, pp. 115–116) notes the following limitation of defining *kill* as <cause to die>: "If I accidentally drop a banana skin in the staircase of the faculty building, and the dean slips on it in a lethal fall, I have caused the dean to die, but I did not kill him in the usual sense of killing." Likewise, Lakoff (1987, pp. 70–71) observes that while *bachelor* appears to mean <an unmarried man>, this definition would lead one to mistakenly treat as bachelors such people as Tarzan (before he entered

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

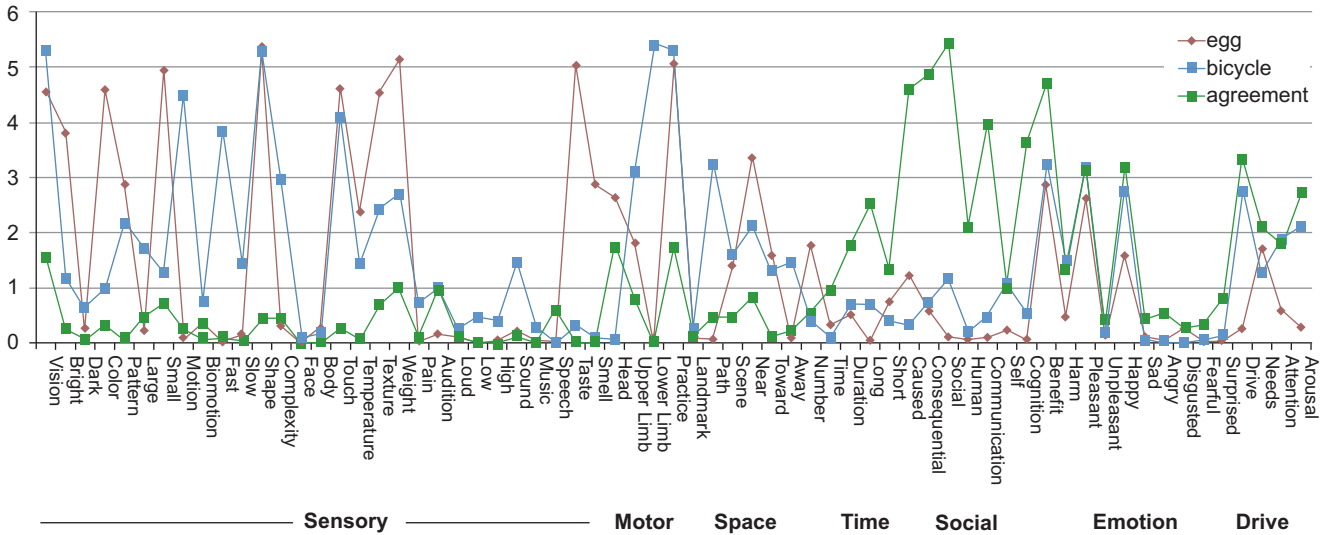


Figure 1.1 Mean ratings on a six-point scale for the words *egg, bicycle*, and *agreement* across 65 experiential attributes that fall into 14 categories (see Table 1.1). (From Binder et al., 2016, p. 161.)

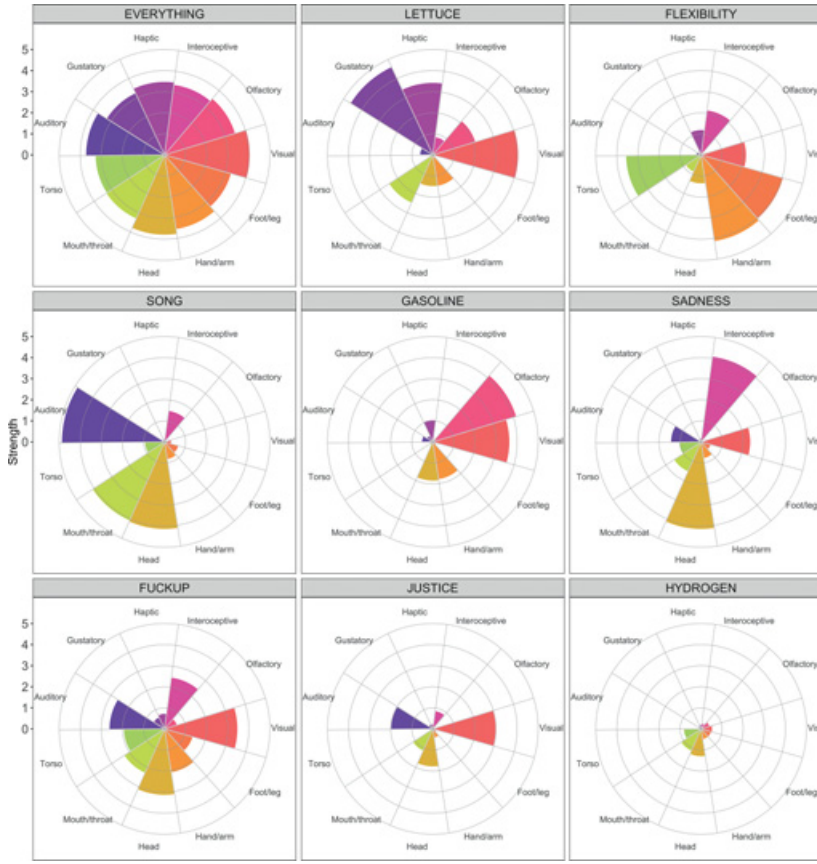


Figure 1.2 Polar plots for nine lexical concepts, showing the mean rating on a five-point scale for each of 11 sensory/motor dimensions. (From Lynott et al., 2020, p. 1278.)

civilization), the Pope, unmarried men with long-term live-in partners, and Muslim men who are permitted four wives but only have three. We'll return to this problem in Section 1.6.

Second, it's not clear when semantic decomposition should stop. As noted above, the meaning of *bird* includes features like <feathers>, <wings>, <beak>, and so on, but these are themselves complex concepts that could be broken down further. How far should such analysis go? Can all word meanings ultimately be characterized in terms of a relatively small pool of irreducible semantic primitives, or what Aitchison (2012, p. 90) playfully calls "primordial atomic globules," and if so, what kind of evidence might support them? Historically, this has been a very contentious topic. However, Binder et al.'s

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

(2016) novel approach may constitute a genuine advance, as their 65 rudimentary types of "experiential attributes" of concepts are based on fundamental brain systems (see again Table 1.2).

Third, mere lists of features cannot fully capture word meanings because they don't indicate how those features form coherent, structured concepts. Continuing with the example of *bird*, its semantic features are not just free-floating, independent properties; instead, they are related to each other in complex ways. A bird's wings, after all, are what enable it to fly. Such relationships between features are essential aspects of word meanings, as discussed in the next section.

1.4 Frames

How a word's semantic features relate to each other is specified by what's often called a frame (Lehrer & Kittay, 1992; Gamerschlag et al., 2014). Just as the frame of a house provides its structural outline, so the frame of a concept establishes its overall configuration. It does this by delineating and classifying the connections between the assorted semantic features, thereby integrating them to form a unified mental representation. Carrying on with the case of *bird*, its complete frame must be quite complex, because appropriate relationships must be captured between myriad features both within and across several semantic spheres, including body plan, behavior, habitat, and relevance to humans. Considering just a bird's body plan, the conceptual frame must specify many structural and functional relationships – for example, the head contains two horizontally aligned eyes, each of which contains an eyeball, which in turn contains a pupil. As with *bird*, other words for objects also tend to have complex frames. For instance, the frame for *car* must link together numerous features involving a car's design, operation, use, and so on. A small fragment of this frame is shown in Figure 1.3.

Verbs likewise have frames that systematically combine multiple semantic features to create cohesive concepts. Because most verbs encode events, their meanings are typically built around a few basic types of features, such as the following:

- entities that play certain roles (e.g., *insert* encodes an agent, an entity that the agent moves, and a partly or fully enclosed space where that entity ends up)
- properties of those entities (e.g., *smear* applies to a semisolid substance)
- manner of motion (e.g., *roll*, *slide*, and *swing* denote different movement patterns)
- path of motion (e.g., *enter*, *exit*, and *ascend* denote different movement directions)

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

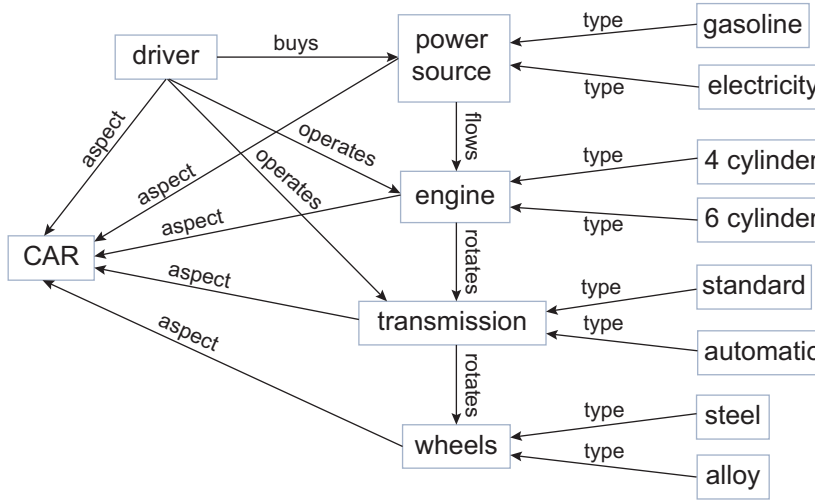


Figure 1.3 A fragment of the semantic frame for *car*. (Inspired by Barsalou, 1992.)

- causation (e.g., *blend*, *braid*, and *carve* encode actions with certain outcomes)
- intention (e.g., *rummage*, *wash*, and *decorate* specify certain agentive goals)

The conceptual frames of verbs organize these and other kinds of features in ways that often involve several hierarchically structured subevents that are related to each other in terms of, say, manner and path of motion, or cause and effect. For example, the meaning of *chase*, as it appears in *The cat chased the mouse*, can be paraphrased roughly as "The cat acts and goes toward the mouse (which is going away from it) in order to be at the mouse" (Pinker, 1989, p. 203). Thus, like the frames for object nouns, those for action verbs can be quite complex.

1.5 Prototypes

Back in the mid-1970s, Eleanor Rosch provided the first psycholinguistic evidence that for many categories of objects people do not treat membership as an all-or-nothing affair based on a set of necessary and sufficient conditions; instead, they regard some objects as being better – that is, more typical or representative – members of a category than others, due to their having more of the expected features (Rosch, 1978).

For instance, in one experiment participants were shown lists of names for objects that belong to certain categories, and for each list their task was to rate

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

on a seven-point scale the extent to which they considered each object to be a "good example" of the given category. Returning once again to the case of *bird*, the results were remarkably consistent: *robin*, *sparrow*, *canary*, *blackbird*, *dove*, and *lark* all came out high; *parrot*, *pheasant*, *albatross*, *toucan*, and *owl* were somewhat lower; *flamingo*, *duck*, and *peacock* were lower still; *ostrich*, *emu*, and *penguin* were more than halfway down the scale; and the lowest rating of all was given to *bat*, which, strictly speaking, doesn't even refer to a bird. In short, some birds are reliably regarded as being birdier than others. What's more, people gave similarly graded membership judgments for other categories of objects, including the following (for each of these categories I mention a central member, an intermediate one, and a peripheral one): *fruit* (*apple*, *watermelon*, *coconut*); *vegetable* (*carrot*, *potato*, *rice*); *carpenter's tool* (*hammer*, *awl*, *axe*); *clothing* (*pants*, *shoes*, *bracelet*); *vehicle* (*automobile*, *tractor*, *elevator*); *weapon* (*gun*, *bow*, *rope*).

Further evidence that conceptual categories have graded structure comes from other experiments. When participants are asked to verify sentences asserting that certain objects belong to certain categories, their response times are faster for central members (e.g., *A robin is a bird*) than for peripheral ones (e.g., *A penguin is a bird*). Also, participants are more willing to accept sentences about the features of objects if those objects are central members of the relevant category (e.g., *A robin can fly*) than if they are peripheral members (e.g., *A penguin can fly*). In addition, during language development, children usually learn words for central category members before peripheral ones. Finally, findings similar to those just mentioned have also been reported for a variety of nonobject categories, such as actions, spatial relations, and abstract notions.

This line of work suggests that categorization is often guided by a mental representation of a prototype. The precise nature of prototypes has been controversial (Hampton, 2006), but for present purposes the following account should suffice. Notably, a prototype is *not* a memory of a single ideal instance of a category, like a particular robin that one saw; in fact, exemplar representations may play only a limited role in knowledge of word meaning (Box 1.1). Rather, a prototype is a continually updated average that, based on the probabilistic distribution of features across category members that one encounters, combines the most common features to form a cognitive reference point. Categorization then involves determining how well the features of a given object (or other stimulus) match those of the current prototype. Thus, the reason why people judge *robin* and *sparrow* to be better examples of *bird* than *penguin* and *ostrich* is primarily because the former share more features with the prototype (e.g., <flies> and <sings>) than the latter.

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

Box 1.1 Is There an Exemplar Theory of Word Meaning?

In the field of category learning, some researchers have argued that summary descriptions like prototypes are not needed to represent categories because memories of particular exemplars are sufficient (e.g., Medin & Schaffer, 1978). The strongest evidence for this view comes from classification experiments in which a small number of artificial exemplars are presented to participants over and over again, and the novel categories they belong to are poorly organized and very hard to learn. Under such circumstances, it makes sense that participants would remember individual exemplars and use them to facilitate classification.

But what about word meanings? Do exemplar representations contribute to them too? Perhaps to some degree in some cases. For instance, my understanding of the word *dog* is probably influenced by idiosyncratic aspects of my own dog Mishka (a large golden doodle) as well as several other dogs that I know well.

On the other hand, Murphy (2016) argues that many properties of word meaning have no clear or plausible explanation in terms of exemplars. One concern involves the high-level forms of knowledge that, as discussed in Section 1.6, grant coherence to large conceptual domains. For instance, people know that, as a rule, all animals are born, grow, eat, breathe, reproduce, and die. The exemplar theory would require these beliefs to be redundantly associated with mental representations of unique exemplars of every animal species. Surely, however, it would be more efficient to store such beliefs as broad generalizations that cover the entire domain? A related challenge is that, as discussed in Section 1.8, object concepts tend to be organized in taxonomic hierarchies. Murphy (2016, p. 1037) points out, however, that according to the exemplar theory, "every squirrel I see would have to be remembered as a squirrel, rodent, mammal, vertebrate, and animal. And I see a *lot* of squirrels." Again, this would be very inefficient, especially considering that such hierarchical knowledge could easily be linked with a squirrel prototype.

A caveat, however, is that recent research indicates that typicality ratings are less consistent than implied by Rosch's seminal studies (Folstein & Dieciuk, 2019). This is largely because people often judge typicality relative to situations that they call to mind, and these situations vary across individuals (Yeh & Barsalou, 2006). For example, chickens might be typical birds on a farm, whereas penguins might be typical birds on a glacier, and each will be more easily processed in the appropriate context.

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

1.6 Mental Models

For many word meanings, the prototype not only captures the most common properties of the category but also reflects a more abstract mental model of the relevant domain of experience. More specifically, a mental model is a structured set of schematic beliefs and expectations about the world, derived from personal observation, cultural norms, and, perhaps in some cases, innate knowledge.

Some mental models have a very broad application. A good example is the one that is frequently called intuitive or folk biology, because it covers all living things. People everywhere, regardless of whether they have any scientific knowledge of genetics, are strongly inclined to conceptualize every plant and animal species as having a hidden essence that determines how that type of organism looks, grows, and reproduces (Atran, 1990). This is what lies behind people's understanding that a caterpillar and a butterfly are actually the same kind of creature, despite having very different forms, behaviors, and habitats. It also drives the insightful judgments that preschoolers make about many aspects of living things, such as that a raccoon with a white stripe painted down its back is still a raccoon and not a skunk, but a dog whose insides have been removed is not really a dog anymore and can no longer perform normal canine actions. Indeed, because parents hardly ever talk about the underlying identities of plants and animals, it's possible that children's predisposition to attribute hidden essences to them reflects an innate component of human biological cognition, one that drives learning and reasoning about natural kinds (Gelman, 2003). On the other hand, some researchers maintain that the vitalistic idea of a life force, which is the heart of intuitive biology, is not innate but instead constructed during cognitive development (Xu, 2019). In any event, what matters here is simply that the meanings of individual words for plants and animals (e.g., *banana*, *carrot*, *bird*, *dog*) not only share but are significantly influenced by a deep-seated mental model that helps people make sense of living things.

Other mental models have a much narrower scope but still shape word meanings in important ways. Consider, for example, the word *mother*. It encodes a complex concept that actually involves a cluster of several mental models that usually coincide but sometimes diverge, due to complications resulting from social practices like adoption as well as medical interventions like egg and embryo implantation (Lakoff, 1987):

- The birth model: The woman who gives birth to a child is its mother.
- The nurturance model: The woman who raises a child is its mother.
- The marital model: The woman who is married to a child's father is its mother.

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

- The genetic model: The woman who contributes to a child's genetic material is its mother.

When these models don't align, confusion and disagreement can emerge over who a person's "real" mother is. And the fact that such uncertainties are often hard to resolve suggests that there is no conventionally dominant model for motherhood. Instead, people must resort to more specialized terms like *stepmother*, *surrogate mother*, and *adoptive mother*.

It's also noteworthy that some mental models incorporate scripts, which represent predictable sequences of events. Take the word *wedding*. It invokes a ritualistic process in which two people (typically a man and woman but sometimes same-sex partners) are married. A minimal American wedding script is centered around the exchange of vows, but a more expansive one includes a number of additional events that unfold in a certain order. For instance, according to one tradition, the ceremony itself is followed by a reception, which in turn is followed by a send-off, and each of these events consists of several subevents. A bit of reflection reveals that many other words, such as *graduation*, *picnic*, and *restaurant*, are also associated with scripts. As Aitchison (2012, pp. 88–89) points out, "anyone going to a restaurant will expect to be led to a table, to be presented with a menu, to order food, to eat that food, then pay for it."

1.7 Background Situations

According to research inspired by the field of situated cognition (Robbins & Aydede, 2009), another important property of many words is that their meanings are stored in long-term memory together with representations of associated background situations (Barsalou, 2009, 2016b; Yeh & Barsalou, 2006). Consider, for example, the word *chair*. Strictly speaking, it only denotes a type of object that affords sitting. But people also know that different kinds of chairs – with varying specifications for size, shape, material, flexibility, comfort, and purpose – tend to be found in different situations. Living rooms tend to contain large stuffed chairs; dentists' exam rooms tend to contain fully reclining chairs; offices tend to contain swivel chairs on wheels; and other settings, like classrooms, movie theaters, passenger jets, and yard parties, tend to contain still other kinds of chairs.

Several studies support the psychological reality of such close couplings between word meanings and background situations. For instance, words for objects associated with a high location, like *bird*, are processed more efficiently when presented at the top versus the bottom of a computer display, and words for objects associated with a low location, like *snake*, are processed more

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

efficiently when presented at the bottom versus the top of a computer display (Šetić & Domijan, 2007; Estes & Barsalou, 2018). In addition, objects are named more accurately when presented in consistent versus inconsistent scenes, and scenes are named more accurately when they contain consistent versus inconsistent objects in the foreground (Lauer et al., 2018; Leroy et al., 2020).

1.8 Word Associations and Co-occurrence Patterns

Another aspect of lexical semantics involves the links *between* words that allow them to form associative networks. Such networks help people recognize and recall words (Locker et al., 2003; Steyvers et al., 2004). In addition, they affect people's judgments about the meanings of words (Solomon & Barsalou, 2004), and there's increasing evidence that they even contribute directly to those meanings (Connell, 2019; Günther et al., 2019; Lenci, 2018; Lupyan & Lewis, 2019; Wingfield & Connell, 2022). Here, we'll focus on the following kinds of lexical links: traditional semantic relations, thematic relations, and collocational relations.

Two types of traditional semantic relations are synonymy and antonymy. Synonymy involves sameness of meaning. However, if two words can only be treated as having the same meaning if they are equally interchangeable in all contexts, then absolute synonymy is quite rare. Near synonymy, though, is widely attested. For instance, *calm* and *placid* have very similar meanings but differ in subtle ways, because it would be fine to say *She was calm just a moment ago* but odd to say *She was placid just a moment ago*. Some other examples of nearly synonymous word pairs are *brave:courageous*, *big:large*, *almost:nearly*, *fog:mist*, and *amble:stroll*. In contrast to synonymy, antonymy involves oppositeness of meaning. Some antonyms are mutually exclusive, like *male:female* and *alive:dead*, but most are defined with reference to a gradable scale, as exemplified by *long:short*, *thick:thin*, *fast:slow*, *hot:cold*, and *heavy:light*.

Several other types of traditional semantic relations derive from taxonomic hierarchies. Such hierarchies have been applied to many conceptual domains, but they are especially useful for representing nested categories of plants and animals. For instance, the domain of plants includes the category of trees, which includes the more specific category of oaks, which includes the even more specific category of white oaks. Within this taxonomy, *tree* is a hyperonym or superordinate of *oak*, and conversely *oak* is a hyponym or subordinate of *tree*. Furthermore, words at the same taxonomic level as *oak*, like *sycamore*, *elm*, and *maple*, are all cohyponyms or coordinates. The same types of taxonomically based semantic relations can also be found in other

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

conceptual domains. For example, in the domain of action, *cook* is a hyperonym or superordinate of *bake*, *bake* is a hyponym or subordinate of *cook*, and *bake*, *roast*, *fry*, and *simmer* are all cohyponyms or coordinates.

Turning to thematic relations, as the name suggests, they include many types of topical or contextual connections between words, such as through objects, events, functions, locations, social organizations, evaluations, emotions, and so on. Thematic relations are usually manifested when people are given a list of cue words and must produce, for each one, the first associated word that comes to mind. To be sure, in such circumstances the links between cues and associates sometimes involve the sorts of traditional semantic relations described above. In most cases, however, they also involve various kinds of thematic relations. For example, when 114 participants were given the cue word *dinner*, their responses were, in descending order of frequency, *supper* (61), *eat* (13), *lunch* (11), *food* (10), *meal* (10), *bell* (1), *dance* (1), *guest* (1), *hungry* (1), *jacket* (1), *party* (1), *repast* (1), *steak* (1), and *table* (1) (Nelson et al., 2004). Some of the most common responses were linked with the cue through traditional semantic relations: *supper* is a near-synonym of *dinner*; *lunch* is a cohyponym; and *meal* is a hyperonym. All the other responses, however, were linked with the cue through diverse thematic relations, such as actions (*eat*), objects (*food*, *repast*, *steak*, *jacket*, *bell*), settings (*table*), social roles (*guest*), social situations and events (*party*, *dance*), and introspective states (*hungry*).

Currently, the richest information about thematic relations comes from the "Small World of Words" project (De Deyne et al., 2019). Word associations were collected for over 12,000 cue words, and in contrast to previous studies, participants responded to each cue by giving not just the first word that came to mind, but the first three. This allowed researchers to discover weaker associations and infer denser networks. For example, Figure 1.4 portrays several interconnected word webs anchored by the abstract term *probability*. The word association norms derived from such data can shed new light on the organization of lexical knowledge. For instance, the observed patterns of word association are highly predictive of people's judgments of semantic similarity.

The last kind of lexical link to be considered here involves collocation, which is defined as "a lexical relation between two or more words which have a tendency to co-occur within a few words of each other in running text" (Stubbs, 2002, p. 24). Collocational relations are also referred to as the distributional patterns of words. As mentioned in the Preface, the first scholar to devote serious attention to them was the mid-twentieth-century linguist J. R. Firth, who famously stated that "you shall know a word by the company it keeps" (Firth, 1957, p. 11). What this dictum implies is that much of one's

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

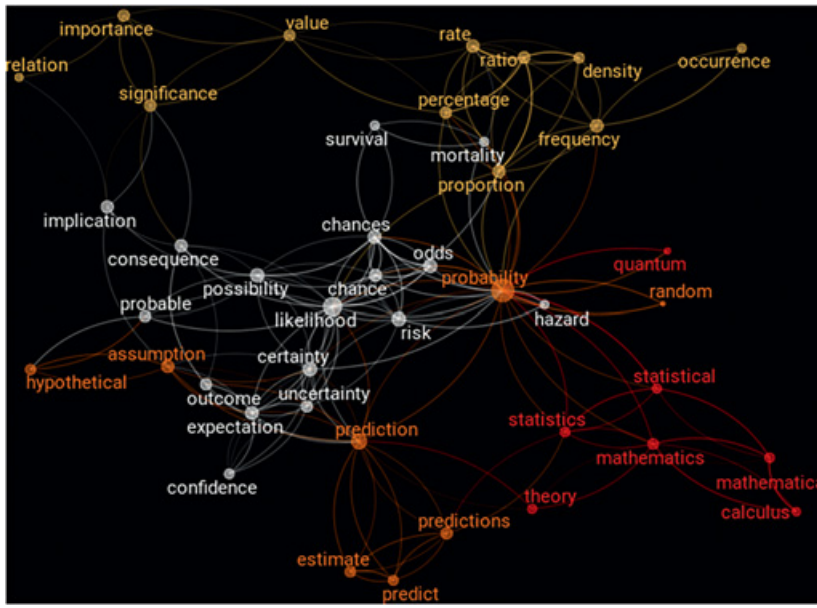


Figure 1.4 Some networks of word associations anchored by the abstract term *probability*. (Image kindly provided by Simon de Deyne.)

knowledge of a word like *democracy* comes from its tendency to occur together with related words like *government*, *country*, *election*, *vote*, *citizen*, *human rights*, and so on. Firth lacked the resources to rigorously test this idea. During the past 30 years, however, it has become increasingly clear that computational systems called distributional semantic models (DSMs) can learn a great deal about the meanings of words based solely on how their phonological/orthographic forms tend to be arranged relative to each other in discourse. In fact, although DSMs vary in their cognitive plausibility, many of them can approximate human performance on an impressive array of tasks, including sentence completions, similarity ratings, typicality ratings, ambiguity resolution, and the comprehension of gist (Günther et al., 2019; Lenci, 2018; Lupyan & Lewis, 2019; Wingfield et al., 2022). Moreover, the most advanced DSMs are now so sophisticated that they can power some of the artificial intelligence (AI) systems that have recently begun to permeate society. Basically, the secret to these systems' success is that, as Lupyan et al. (2020, p. 938) put it, "the co-occurrence statistics between words can act as a kind of echo of real-world linkages and causal relationships." We'll look more closely at how DSMs work in Chapter 5.

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

1.9 Contextual Variability of Word Meaning

"Word meanings cannot be pinned down, as if they were dead insects. Instead, they flutter around elusively like live butterflies." This picturesque passage from Aitchison (2012, p. 54) reflects the fact that many words have multiple interpretations that are highly sensitive to context. The extreme form of such semantic variability can be found in figurative language, as illustrated by the playful use of the word *lunch* in the following sentence: *"This is my lunch," said John, waving a five-pound note* (Cruse, 2011, p. 108). Here, however, we are concerned with more conventional cases in which a single word form, taken by itself, is ambiguous because it has, or appears to have, two or more established meanings, these being the sorts of meanings that a lexicographer would ideally include in a dictionary. In such cases, different interpretations are appropriate in different contexts.

Let's begin by distinguishing between two types of ambiguity: homonymy and polysemy. Homonymy involves a single word form with two or more seemingly unrelated meanings, as exemplified by the noun *bank*, which denotes either the place where you put your money or the land along the edge of a river. In cases like this, we are dealing with two separate words, each of which has its own meaning. In contrast, polysemy involves a single word form that appears to have two or more senses that are intelligibly connected, as exemplified by *window*, which denotes either the glass (*He shattered the window*), the frame (*He painted the window*), the moveable structure (*He shut the window*), or the gap in a wall (*He leaned out the window*). In cases like this, we are dealing with a single word that arguably has several senses.

Polysemy is a complex topic that has been intensively studied and debated (Cuyckens & Zawada, 1997; Falkum & Vicente, 2015; Gries, 2015; Nehrlich et al., 2003). Perhaps the most controversial issue is this: If one suspects that a word is polysemous, how many different senses should one posit? Consider, for instance, the verb *open*, which applies to a broad range of entities, including eyes, hands, wounds, doors, windows, boxes, books, envelopes, stores, computer files, financial accounts, and human minds, to name but a few. How many distinct but related senses does this word have? Not surprisingly, dictionaries are inconsistent on this score. As Taylor (2012, pp. 220–221) points out, "the *Collins English Dictionary* (1979 edition) distinguishes fourteen different senses, the *Collins Cobuild* (1987) lists twenty-nine, while the *Longman Dictionary of Contemporary English* (1978) identifies only six."

Somewhat more systematic results have been obtained by linguists. One important discovery is that a variety of regular patterns of polysemy seem to

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

generalize across certain sets of words (Pustejovsky, 1995). Some examples provided by Geeraerts (2010, p. 148) include "the relationship between product and producer (*The journal lies on the table/The journal was taken over by Rupert Murdoch*), between process and result (*My purchases took me under an hour/The purchases are still in the trunk of the car*), or between contents and container (*The whole lecture hall laughed/The lecture hall is at the end of the corridor*)." Additional patterns involve the associated senses of words that can be used as either nouns or verbs, like *hammer*, *talk*, and *father*. Furthermore, research on language change has identified many widely attested historical processes whereby nouns and verbs give rise to grammaticalized items with derived meanings (Heine & Kuteva, 2002). For instance, nouns for body parts are often the source of prepositions for spatial locations, as when a noun for <head> evolves into a preposition for <on top of>, or a noun for <belly> evolves into a preposition for <in>, or a noun for <butt> evolves into a preposition for <behind>. In fact, this last path is what yielded the polysemy of the English word *behind*.

On the other hand, when the focus shifts to less regular cases of apparent polysemy, linguists disagree strongly over the proper analytic approach. Some are quite happy to posit myriad, minimally different senses of words (e.g., Lakoff, 1987). From the perspective of online sentence comprehension, an advantage of this "splitting" approach is that it allows listeners and readers to select from a large pool of possible word senses the one that best fits the given context, but a shortcoming is that this selection process is likely to place heavy demands on executive function, as rampant ambiguities must be quickly resolved. At the opposite extreme, other linguists are reluctant to ever posit more than one sense of a word (e.g., Ruhl, 1989). This "lumping" approach has the virtue of significantly reducing lexical–semantic ambiguity, but it faces the daunting if not frequently insurmountable challenge of identifying unitary meanings for words that are used in quite heterogeneous contexts, like *open*. Charting a middle course, yet another strategy is to freely posit multiple senses of words, but only when necessary, which is to say, when the proper interpretation of a particular nonfigurative use of a word cannot be determined from previously recognized senses, contextual cues, or general knowledge. Adopting this "principled polysemy" approach, Evans and Green (2006) posit 15 distinct senses of the preposition *over* (Table 1.3). This may seem like a lot of senses, but it's actually far fewer than what other researchers have ascribed to this semantically complex word (for a brief summary of "the story of *over*" see Taylor, 2012, pp. 233–238).

Large-scale corpus-based analyses, including DSMs, can facilitate our understanding of polysemy, not only by revealing rare uses of certain words (Gries, 2006), but also by allowing researchers to measure the semantic

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

Table 1.3 An illustration of polysemy involving the related yet distinct senses of *over*. (From Evans & Green, 2006, p. 348.)

	Sense	Example
1	ABOVE (CENTRAL SENSE)	*The picture is over the sofa*
2.A	ON-THE-OTHER-SIDE-OF	*St. Paul's is over the river from Southwark*
2.B	ABOVE-AND-BEYOND (excess I)	*Your article is over the page limit*
2.C	COMPLETION	*The movie is over*
2.D	TRANSFER	*The discredited government handed over power*
2.E	TEMPORAL	*The relationship had altered over the years*
3	COVERING	*The clouds are over the sun*
4.A	EXAMINING	*Mary looked over the document quite carefully*
4.B	FOCUS-OF-ATTENTION	*The committee agonized over the decision*
5.A	MORE	*Jerome found over 40 kinds of shells on the beach*
5.A. 1	OVER-AND-ABOVE (excess II)	*The heavy rains caused the river to flow over its banks*
5.B	CONTROL	*She has a strange power over me*
5.C	PREFERENCE	*I would prefer tea over coffee*
6.A	REFLEXIVE	*The fence fell over*
6.B	REPETITION	*After the false start, they started the race over*

diversity of contexts in which certain words occur (Hoffman et al., 2011). Nonetheless, it remains notoriously difficult to find out how many different senses of putatively polysemous words speakers actually represent in their heads (Sandra & Rice, 1995). After all, as Taylor (2012, p. 227) points out, "all that is needed of the speaker is that she is familiar with the word combinations and the kinds of situations they refer to."

In this connection, it's worth considering how the notion of mental models (discussed in Section 1.6) might simplify some of the problems posed by polysemy. In an insightful passage, Taylor (2012, p. 225) draws on this notion to suggest that the word we used to illustrate polysemy at the outset of this section, namely *window*, might not really be polysemous: "For *window*, we might argue that the unitary meaning refers to our conception of a particular kind of architectural feature. Our concept includes knowledge of the function and usual location of windows, their structure, the materials they are made from, and how we typically interact with them. Any specific use of the word is likely to highlight only some facet of the broader concept . . . ."

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

1.10 Cross-Linguistic Similarities and Differences

Finally, it's important to note that nearly 6,500 mutually unintelligible languages are spoken on our planet (Hammarström, 2016), and as Pagel (2012, p. viii) points out, "this is more different systems of communication in a single mammal species – for that is what we are – than there are mammal species." Within the broad discipline of linguistics, the division that concentrates on cross-linguistic similarities and differences is called typology, and within it the branch that focuses even more narrowly on conceptual representation is called semantic typology (Evans, 2011, 2022; Majid, 2015; Malt & Majid, 2013; Moore et al., 2015). Although this area of specialization is still maturing, it has begun to flourish during the past few decades, as shown by increasingly systematic and insightful studies of numerous conceptual fields, including colors, plants, and animals (i.e., ethnobiology), body parts, landscapes, temperatures, odors, spatial relations, several spheres of action (e.g., locomotion, giving, cutting and breaking, putting and taking, and eating and drinking), and many domains typically encoded by grammatical morphemes and constructions (e.g., number, gender, person, possession, tense, aspect, mood, and evidentiality).

These investigations have shown that languages all over the world display many intriguing similarities in both the content and structure of the concepts that they conventionally express. These cross-linguistic similarities may stem from commonalities in people's bodies, brains, communicative behaviors, and physical and social environments. At the same time, however, research in semantic typology has also revealed a tremendous amount of cross-linguistic diversity in the lexical partitioning and packaging of particular conceptual fields. Because these findings have not yet had much impact on cognitive neuroscientists who study word meaning, we will not consider them in this book. Nonetheless, I encourage readers to keep this topic in mind, as it warrants greater attention from a neuroscientific perspective (Kemmerer, 2019, 2023a, 2023b).

1.11 Summary

- Word meanings are publicly shared concepts for social coordination. Because all competent speakers of a given language have essentially the same basic vocabulary, they can use it to exchange information about all kinds of topics for all kinds of reasons.
- Word meanings consist of features. For example, the meaning of *bird* can be broken down into numerous concrete features, like <feathers>, <wings>, and <beak>, each of which can be decomposed further.
- Frames specify how the features that make up word meanings are organized to form coherent concepts. For example, the frame for *bird* specifies how its

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

myriad features are related to each other both within and across semantic domains such as body plan, behavior, habitat, and relevance to humans.

- Prototypes pull together the most frequently encountered features of the kinds of objects, events, and abstract phenomena captured by word meanings. For example, the prototype for *bird* represents common features like <flies> and <sings>.
- Mental models are schematic sets of beliefs and expectations that contribute to multiple word meanings. For example, some general aspects of the meaning of *bird* derive from folk biology, which specifies that every living species has a hidden essence that determines how that type of organism looks, grows, and reproduces.
- Word meanings are often associated with background situations. For example, birds are associated with nests in trees.
- Word meanings are influenced by several kinds of lexical relations, including the following. Taxonomic hierarchies – for example, ospreys are hawks, which are birds, which are animals, which are living things. Thematic relations – for example, when you think about birds, you might call to mind a variety of contextually associated objects and events, like eggs, chicks, flying, and migrating. Collocational relations – for example, the word *bird* tends to be used together with many other topically relevant words, like *wing*, *feather*, *flock*, *wildlife*, *roost*, *nest*, and so on, and there's mounting evidence that the meaning of *bird* can be derived in part from the statistical co-occurrence patterns of these phonological/orthographic forms.
- Many words are polysemous – that is, they have two or more distinguishable but related meanings. For example, while the avian sense of *bird* is its primary meaning, in British English this word can also be used to describe a pretty, unmarried woman, although this is rather old-fashioned slang and is considered offensive by some.
- The research discussed in this book is based almost entirely on English, but it's important to remember that there are nearly 6,500 languages in the world, and although their word meanings exhibit many similarities, they exhibit far more differences.

Further Reading

Aitchison, J. (2012). *Words in the mind* (4th ed.). New York, NY: Wiley-Blackwell.

Barsalou, L. W. (2016). Situated conceptualization: Theory and applications. In Y. Coello & M. H. Fischer (Eds.), *Foundations of embodied cognition*, *Vol. 1*: *Perceptual and emotional embodiment* (pp. 11–37). Chichester: Psychology Press.

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press

Connell, L. (2019). What have labels ever done for us? The linguistic shortcut in conceptual processing. *Language, Cognition, and Neuroscience*, *34*, 1308–1318.
Croft, W., & Cruse, D. A. (2004). *Cognitive linguistics*. Cambridge, UK: Cambridge University Press.
Cruse, A. (2011). *Meaning in language* (3rd ed.). New York, NY: Oxford University Press.
Geeraerts, D. (2010). *Theories of lexical semantics*. New York, NY: Oxford University Press.
Geeraerts, D., & Glynn, D. (Eds.) (in press). *Cambridge handbook of lexical semantics*. New York, NY: Cambridge University Press.
Malt, B. C., & Majid, A. (2013). How thought is mapped into words. *WIREs Cognitive Science*, *4*, 583–597.
Miller, G. A. (1996). *The science of words*. New York, NY: Scientific American Library.
Murphy, G. (2004). *The big book of concepts*. Cambridge, MA: MIT Press.
Reilly, J., Shain, C., Borghesani, V., Kuhnke, P., Vigliocco, G., Peelle, J. E., Mahon, B. Z., Buxbaum, L. J., Majid, A., Brysbaert, M., Borghi, A. M., De Deyne, S., Dove, G., Papeo, L., Pexman, P. M., Poeppel, D., Lupyan, G., Boggio, P., Hickok, G., . . . Vinson, D. (2025). What we mean when we say semantic: A consensus statement on the nomenclature of semantic memory. *Psychonomic Bulletin & Review*, *32*, 243–280.

https://doi.org/10.1017/9781009436328.002 Published online by Cambridge University Press