

# Word Processing

## 3

*Polysynthetic languages go overboard ... packing whole English sentences into a single word, as in Cayuga Eškakheho. na'táyęthwahs, "I will plant potatoes for them again."*

NICHOLAS EVANS AND STEPHEN LEVINSON

This chapter focuses on the mental lexicon, what information it contains, and how that information is accessed and used in real time as people interpret utterances. Big questions about words in language science include: How do we mentally represent word forms? How are those representations organized? How are word meanings represented in the mind? When we hear or see a word, how do we go about searching our memories for a matching form? What parts of the brain are involved in storing and accessing word meanings and what neural events support word processing?

Language consists of two components, a *lexicon* that captures information about words, their components, and their meanings, and a *grammar* that lays out principles governing how words can be combined into phrases and sentences. The distinction between words and longer expressions is not always as neat and tidy as it is in English, an *analytic* language. We can learn important lessons by looking at analytic language properties. Let's put aside for the time being questions relating to other (very interesting) systems.

*Introduction to Psycholinguistics: Understanding Language Science, Second Edition.*  
Matthew J. Traxler.

© 2023 John Wiley & Sons Ltd. Published 2023 by John Wiley & Sons Ltd.  
Companion Website: [www.wiley.com/go/traxler/psycholinguistics2e](http://www.wiley.com/go/traxler/psycholinguistics2e)

The Anatomy of a Word: How We Mentally Represent Word Form

Lexical Semantics

*Associationist accounts of word meaning*

*The symbol-grounding problem*

*Embodied semantics*

Lexical Access

*First-generation models*

*Second-generation models*

*Third-generation models*

Lexical Ambiguity Resolution

*Does context influence meaning selection for ambiguous words?*

The Neural Basis of Lexical Representation and Lexical Access

*How are word meanings represented in the brain?*

Summary and Conclusions

Test Yourself

To understand how words are represented and processed, we need to **subject them** to different kinds of analysis, because we represent information about words in at least two distinct ways. First, we mentally represent the *form* that words take—the way they sound and the way they look. The way they sound is captured by a *phonetic* or *phonological* code, the way they look by an *orthographic* code, and what they mean by a *semantic* code. When we talk about how word representations are organized, we can focus on different kinds of mental representation. Words may be related to one another because they sound similar (gave–cave), because they look similar (wow–mow), or because they have similar meanings (horse–donkey). Word forms are represented in *lexical networks* and word meanings are stored in a separate, but linked, *semantic memory* or *conceptual store*. To understand how words are represented and processed, we have to be clear whether we are talking about *form* or *meaning*, and we have to recognize that the mind represents these attributes in different ways in separate, but linked systems (Balota et al., 2006; Collins and Loftus, 1975; Hutchison, 2003; McClelland and Rumelhart, 1985).

## The Anatomy of a Word: How We Mentally Represent Word Form

Words are made up of parts. In the same way that we can analyze molecules as being made up of different kinds of atoms, and we can analyze atoms as being made up of different kinds of particles, so we can divide words up into their subcomponent parts. Different psychological and linguistic theories emphasize different aspects of words, and different theories make different claims about which parts of words have the biggest impact on mental processes that activate stored information about words. Let's start by reviewing some classic approaches to word representation.

Classic approaches to word form representation view words as involving a hierarchical arrangement of components. In speech, the lowest level of organization is the *phonetic* feature. Phonetic features, like place and manner of articulation, combine to produce the next level of organization, the *phoneme*. Phonemes can be combined to make up *bigrams* (pairs of phonemes) and *trigrams* (triplets), or we can think of combinations of phonemes as composing *syllables*, consonant–vowel (CV) or consonant–vowel–consonant (CVC) combinations. (CV and CVC combinations result from the fact that when we talk, we alternately open and close our jaws, starting and stopping the flow of air—we literally flap our jaws when we speak.) Syllables themselves can be divided into *onsets* (the initial CV combination, like *spa* in *spam*) and *rimes* (the ending VC combination, like *am* in *spam*).

One or more speech sounds can combine to produce a *morpheme*—defined as the smallest unit in a language that can be assigned a meaning. One or more morphemes can be combined to produce a word. *Cat*, for instance, is a *monomorphemic* (“one morpheme”) word because only one morpheme makes up the word. Languages also combine morphemes to produce *polymorphemic* (“more than one morpheme”) words, as in the compound word *blackboard* (some languages, like Turkish, Finnish, and German are prolific combiners of morphemes).<sup>1</sup> Languages also provide ways of changing the flavor of a word meaning. We can alter the meaning of *cat* (a *singular* noun, used to refer to one animal) by adding a *bound morpheme*, *-s*, resulting in the polymorphemic word *cats* (a *plural* noun, used to refer to more than one animal).

Words convey meaning (and more) from speaker to listener. How is that accomplished? To begin, we have to discriminate between two different definitions of the term *meaning*: *sense* and *reference* (Jackendoff, 1983). *Sense* refers to dictionary-like or encyclopedic knowledge that we have about words. So, for example, the word *cat* maps on to generic information about form and function. When we hear *cat* we can access the information that cats are mammals, they have fur, they are kept as pets, and so forth. When we hear *knife*, we think of metal objects used for cutting things. *Reference* is another type of meaning. When we use words to refer to people, objects, or ideas the words themselves have senses, but their specific meaning in a given context depends on what the words point to—what they *refer* to.

Consider the situation in Figure 3.1. This mini-universe contains two objects. If someone wants to direct your attention to one of the objects, they need to craft an utterance that refers to that object. One could use a number of different expressions to point to either object. Each of these different expressions will have a different sense. Let's say a speaker chooses to refer to one of the objects as *The dark orange one*. The sense of the words *dark* and *orange* helps the listener pick out the object on the left. The speaker could have said, *The one on the left*. That expression picks out the same object as *the dark orange one* did, so it has the same referent, and in that way the two expressions “mean” the same thing—both expressions direct your attention toward the same object. But the two expressions have different senses—being dark orange is not the same thing as being on the left—and so the two expressions “mean” different things at the level of sense. Different expressions that have the same sense can have different referents in different contexts. If our speaker said, *The bigger one* in the context of Figure 3.1, that would point to the dark orange object. If they said the same thing in the context of Figure 3.2, that would point to the pale orange object.




Figure 3.1 A two-object universe




Figure 3.2 Another two-object universe

We can talk about the meaning of a word by referring to the sense of the word; and we can talk about the meaning of a word by focusing on what the word refers to. Chapter 6 discusses reference in detail, so this chapter will deal exclusively with the *sense* meaning of words. When this chapter talks about *semantics* or meaning, think “sense.”

How are word meanings (*senses*, that is) represented in the mental lexicon? And what research tools are appropriate for investigating word representations? One approach to investigating word meaning relies on introspection—thinking about word meanings and drawing conclusions from personal experience. It seems plausible, based on introspection, that entries in the mental lexicon are close analogs to dictionary entries. If so, the lexicon (like the dictionary) includes information about words’ functions (what categories do they belong to, *verb*, *noun*, *adjective*, etc.), which determines how they can combine with other words (adverbs go with verbs, adjectives with nouns). This dictionary approach to the lexicon assumes that individual words refer to *types*—i.e. the core meaning of a word is a pointer to an interchangeable set of objects in the world (Gabora et al., 2008). Each individual example within a category is a *token*. For instance, *team* is a type, and *Yankees*, *Twins*, and *Mudhens* are tokens of that type.<sup>2</sup>

If word meanings are types, how do we represent types? We *could* represent a type by making a list of defining, necessary, or core characteristics. Some words seem to be easily represented by a small number of core, necessary features. “Bachelor,” for example, seems to be well represented by combining the concepts “human,” “adult,” “male,” and “unmarried.” However, this apparent simplicity could be misleading.

Consider, for instance, the concept “cat.” We could use its core features (e.g. “cat” = “cute + furry + killing machine”) to represent its meaning. But we know an awful lot more than that about cats (they have claws, they see well at night, they cough up hairballs, they don’t make good doorstops, you can’t use them to iron your clothes, etc.). The question then becomes, of all of the millions of things one could include in the dictionary entry under the word *cat*, which things get put in and which things get left out? Does the meaning of *cat* include the fact that it can breathe? Does it include the fact that it is larger than a tomato and smaller than an airplane? Probably not. But where do you draw the line? Which properties are prestored in long-term memory, and which are derived “on the fly?” What we really need to store to represent the meaning of the word *cat* is just its core or essential properties—those things that make up the essence of “cat” and that differentiate between cats and other things. In which case, we might store just features like “mammal, feline, pet, makes purring sound” and perhaps a visual image of a prototypical cat.

This “core features” approach runs into trouble very quickly, however, as many fairly easy to understand concepts do not have consistent core properties across different versions of the concept. Even apparently simple concepts like “bachelor” run into trouble (Pinker, 1994). Are monks bachelors? Not really, but they certainly are human, adult, male, and unmarried. The concept *game* is fairly common, and different activities are easy to categorize as games, but a single feature or combination of features seems to be consistent across all of the things we identify as being a game (Gabora et al., 2008; Murphy and Medin, 1985). If the concept “game” does not have any necessary or universal features, what do we list as core properties in the lexicon’s entry for the word *game*?

Another issue is that some referents seem to make better examples of a category than others. Most people judge *red hair* as being a worse example of the word *red* than *fire engine red* (Rosch, 1973). If word meanings are based on types made up of fully interchangeable tokens, then every instance of *red* should be just as good as every other instance of *red*. Finally, many categories are a bit “fuzzy” or vague—it is not clear where exactly one category stops and another one begins. If the whole point of categories is to include and exclude tokens on the basis of core features, fuzzy boundaries are problematic.

These are the kinds of problems that have led many language scientists to abandon the “defining” or “core” features approach to lexical semantics. Until someone comes up with a better scheme, dictionary-definition-like entries do not seem to be a good way of explaining how word meanings are represented in the mental lexicon.

One way to sidestep problems associated with the dictionary entry theory of semantics is to operationalize word meanings as reflecting collections of associated concepts. According to this type of account, a word meaning is defined as “whatever comes to mind when someone says the word.” *Semantic network* theory follows this approach (Collins and Loftus, 1975; Collins and Quillian, 1972; see also Rips et al., 1973; Smith et al., 1974) and has been the dominant theory in artificial intelligence approaches to semantics for decades (see Ober and Shenaut, 2006, for a review; related approaches include Ken McRae’s feature-based semantic nets; McRae and Boisvert, 1998; McRae et al., 2005, 1997).

Semantic network theory proposes that a word’s meaning is represented by a set of *nodes* and the *links* between them (as in Figure 3.3). The nodes represent concepts whose meaning the network is trying to capture, and the links represent relationships between concepts. For example, the concept *goose* would be represented as an address in memory (a node) connected to other addresses in memory by different kinds of links. One of the important kinds of links in semantic network theory is the “*is a*” type. The *is a* link encodes relationships between general categories and the concepts that fall within the category. So, *goose* would be connected to the *waterfowl* node with a unidirectional *is a* link (representing the concept that *a goose is a waterfowl*). The *waterfowl* category node could be connected to many different instances (*duck, goose, coot, swan, seagull*, and so forth), and could in turn be connected to a superordinate category node, like *bird*, with yet another *is a* link. Subordinate concepts, like *goose*, inherit the properties of superordinate nodes via transitive inference (a goose is a waterfowl, a waterfowl is a bird, therefore a goose is a bird). There is no need to directly connect the specific concept *goose* to the more general concept *bird*, and this helps conserve memory resources. The meaning of a word is represented as the pattern of activation in the network triggered when the corresponding node is activated.

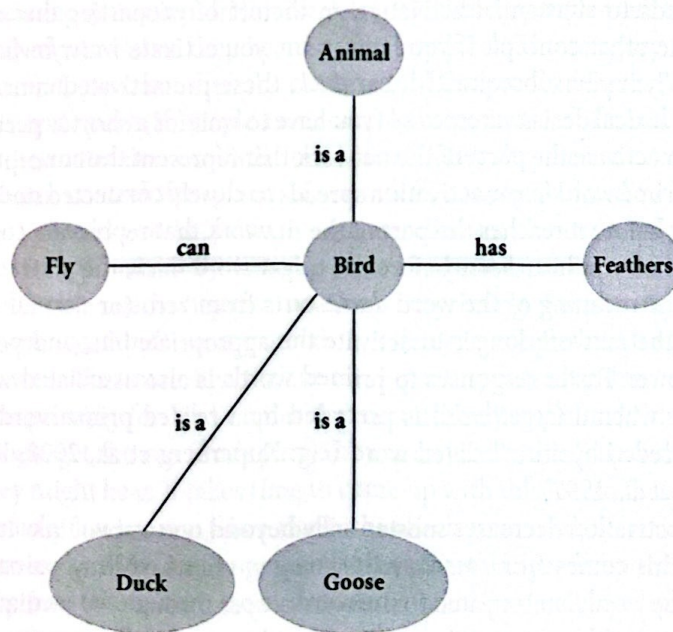


Figure 3.3 A piece of a semantic network

In early work, Collins and Quillian showed that statements such as *A canary can fly* primed responses to statements such as *A canary is a bird*. The explanation for this effect was that reading *A canary can fly* caused activation to spread from *canary is a bird* to *a bird can fly*. So hearing *A canary can fly* entails implicitly activating the relationship *a canary is a bird*, and that property is already activated when subjects read *a canary is a bird* (Collins and Quillian, 1970). Other kinds of nodes and links are used to represent other properties and attributes of individual concepts, like *goose*. For example, *has* links and *can* links connect concepts to components (a *goose has feathers, a beak, and wings; a goose can fly*).

The idea of *spreading activation* is used to explain how information represented in the semantic network is accessed, and why words that are related to one another facilitate access to one another (Collins and Loftus, 1975; Posner and Snyder, 1975). Spreading activation is a hypothetical mental process that takes place when one of the nodes in the semantic network is activated. If someone says *goose*, the *goose* node is activated by the matching phonological (sound) or orthographic (spelling) information. Activation from the *goose* node then spreads to nodes that are connected to it. Activating the *goose* node causes activation to spread to the superordinate node, *bird*, and to the attributes connected to *bird*, *has wings*, *has feathers*, and *can fly*. Spreading activation has two important properties: (a) it is automatic. It happens very fast and we cannot control it; (b) it diminishes the further it has to go. Like ripples in a pond, nodes that are directly connected to *goose* are strongly and quickly activated when you see or hear *goose*. More distantly connected nodes are less strongly and less quickly activated; beyond a couple of degrees of separation, no changes in activation should occur.

The properties of spreading activation help explain how people respond during *priming* tasks. *Priming* occurs when presenting one stimulus at time 1 speeds the response to another stimulus at time 2. In classic work on word processing, people respond faster in *lexical decision* and *naming* experiments when a target word like *duck* is preceded by a related word like *goose*, compared to a control condition where *duck* is preceded by an unrelated word like *horse* (Meyer and Schvaneveldt, 1971, 1976; Moss et al., 1995). This kind of priming is referred to as *semantic priming*. Semantic network theory explains semantic priming as resulting from the spread of activation in the semantic network. Because *duck* and *goose* have many attributes in common, activating one of the concepts necessarily leads to substantial activation in the set of properties that makes up the meaning of the other concept. If you hear *goose*, you activate *waterfowl*, *bird*, *feathers*, and *can fly*. When you subsequently hear *duck*, those pre-activated concepts speed up the naming or lexical decision response (you have to wait for a shorter period of time for the network to activate the parts of the network that represent the concept *duck*). When you hear the prime word *horse*, activation spreads to closely connected nodes, but activation dies away before it reaches the part of the network that represents concepts related to *duck*. So, when you hear *horse* before the target word *duck*, the pattern of activation representing the meaning of the word *duck* starts from zero (or *normal resting activation*), it takes the network longer to activate the appropriate bits, and your behavioral response is slower. Faster responses to primed words is also associated with decreased neural activity when a target word is preceded by a related prime word compared to when it is preceded by an unrelated word (e.g. Kuperberg et al., 2008; Rissman et al., 2003; Wagner et al., 1997).

Spreading activation decreases substantially beyond one or two links in the network. Evidence for this comes from *mediated priming* studies involving pairs of words like *lion-stripes*. The word *lion* is related to the word *stripes* through the mediating word *tiger* (*lion* is associated with *tiger*, *tiger* is associated with *stripes*). When you hear *lion*, activation spreads to *tiger*. When *tiger* gets activated, it should cause activation to spread to

A number of experimental tasks are used to investigate word processing. Two of the most common ones are *lexical decision* and *naming*. In the lexical decision task, people are presented with lists of stimuli, either auditorily or visually on a computer screen. Some of the stimuli are real words, like *cat*, *dog*, *bachelor*, and some are not, like *wat*, *rog*, and *lachenor*. The individual's task is to indicate, as quickly as possible, whether the stimulus is a word or not. If you have an entry in your mental dictionary that corresponds to the stimulus,

you will say "yes," otherwise you will say "no," and the amount of time it takes you to respond is an index of how easy it was to access the word's entry in the lexicon. *Naming* also (usually) involves lists of words, and you respond by saying the word out loud as quickly as you can, but it is not necessary to present nonwords. Here again, the amount of time it takes you to say the target word indicates how long it takes you to access the lexicon and find the word you are trying to say (see Balota et al., 2006; Potts et al., 1988).

*stripes*. If so, then *lion* should prime your response to the word *stripes*. In fact, hearing or reading the word *lion* does lead to a small priming effect for the word *stripes*, so activation does spread beyond directly connected concepts. But if activation can spread beyond immediately connected nodes, what prevents activation from spreading all over the network? If it did, everything in the network would be activated every time you heard any word.<sup>3</sup> According to semantic network theory, the total amount of activation that can be spread is limited. Consequently, nodes directly connected to the prime word are strongly activated, but less directly connected nodes are less strongly activated, with activation decreasing with increased distance in the network. And, in fact, *lion* primes *stripes* much less than it primes the directly related word *tiger* (Balota and Lorch, 1986; De Groot, 1983; Kumar, 2021; McNamara and Altarriba, 1988).

Behavioral evidence suggests that word-to-word associations are activated quickly, without conscious effort, and outside of our control, supporting the automatic part of automatic spreading activation. In Jim Neely's (1977) study, people were told to expect a particular kind of word after they heard a category label. The category label might be something like *body part*, but the subjects were told that words referring to birds would follow the cue *body part*. If people can control the activation of concepts, then they should focus their attention on birds immediately after they hear or see the cue *body part*. If people can control the spread of activation, concepts related to *body part* (like *arm*, *leg*, *hand*) should not be primed, and members of the expected category *birds* should be primed. When Neely tested people's responses to expected (bird) targets immediately after the cue *body part*, there was no priming. But the unexpected, body part names (*arm*, *leg*, *hand*) were primed. If a delay (a couple hundred milliseconds) intervened between the time when the cue (*body part*) appeared and the expected target appeared, then priming for bird names did occur. Neely explained this by proposing that, when people get the cue (*body part*) they strategically think up a short list of bird names that they might hear. It takes time to come up with this list, so there is no priming for birds right away, but later on when the list has been generated, there is a good chance that the target word will be on the generated list, and this speeds up the response. The pattern of response (immediate, fast reaction to body parts; delayed priming of bird names) is consistent with two processes: fast, automatic activation spreading from the cue to related concepts and slower, non-automatic (strategic) attention shift to a

self-generated list of bird names. The existence of fast, automatic spreading activation and a slower strategic modulation of word activation levels is also supported by data showing that some aphasic patients appear to have intact automatic priming, but impaired strategic priming. In experiments with short lags between primes and targets, these aphasic patients show normal levels of priming, but at longer lags, no priming is observed (Hagoort, 1997; Ostrin and Tyler, 1993).

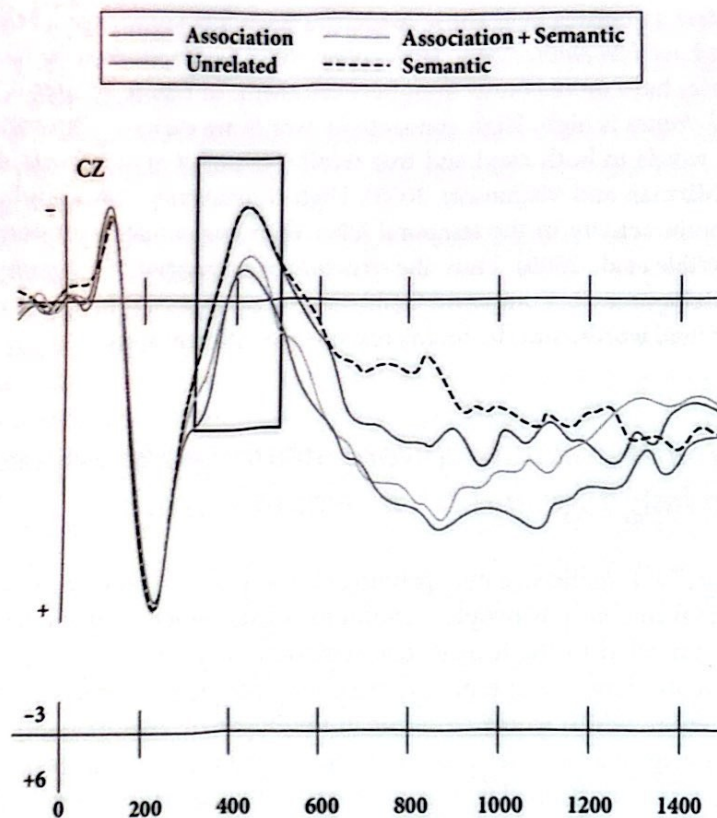
According to semantic network theory, words are related to one another because they have links to shared nodes. *Duck* and *goose* both connect to the *bird* node, the *feathers* node, and so forth. Two words can prime one another because they have similar representations due to shared nodes. This leads to the kinds of priming effects described earlier and also influences what happens to semantic knowledge when the brain is damaged (e.g. Lampe et al., 2021; Moss et al., 1998). Two words can also be related to one another, whether they share nodes or not, if the two words co-occur in the language. *Police* and *jail* can prime one another, not because police officers resemble jails or vice versa, but because the two words appear together often, so the presence of one of the pair may be used to predict the imminent appearance of the other (as in classical conditioning theory; Skinner, 1957). One of the challenges in word-processing research is determining whether priming effects (like *duck-goose* priming) result from sharing nodes in a network, which is the classical view of semantic priming, or whether priming occurs simply because words co-occur.

Although the degree of priming in an experiment depends critically on what kinds of tasks are used and how stimuli are displayed, robust priming is observed for pairs of words that are associated with one another (Moss et al., 1995; Perea and Gotor, 1997; Shelton and Martin, 1992). Priming is harder to detect when pairs of words share elements of meaning, but are not associated, especially when the semantic relationship consists of belonging to the same general category (like *animal* or *clothing*). So, although *pig* and *horse* come from the same category (*animal*, or even more specifically *farm animal*), the priming between *horse* and *pig* is more fragile than between pairs of words that have an associative relationship (like *dog* and *cat*).<sup>4</sup>

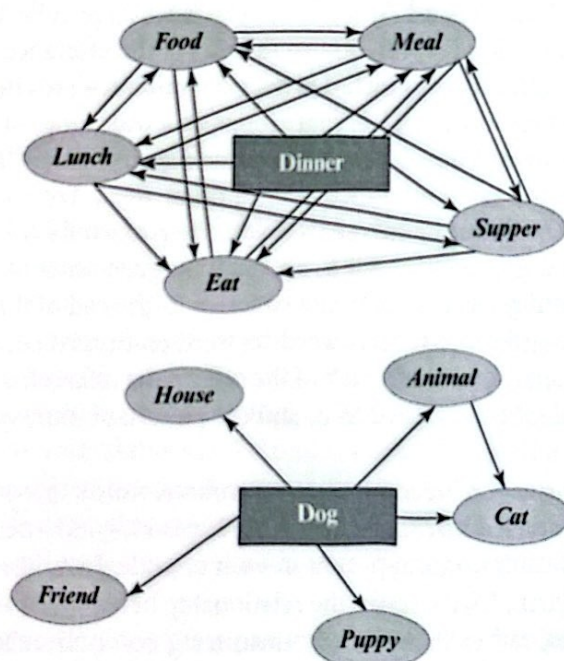
We do not know for sure whether purely semantic relationships (*horse-pig*) produce priming effects in tasks that tap automatic meaning activation, but there is a growing consensus that associative priming and semantic priming are governed by different mechanisms. For instance, Sinéad Rhodes and David Donaldson (2008; see also Farshad et al., 2021) conducted an event-related potential (ERP) experiment where they showed subjects pairs of words that were only associatively related (*fountain-pen*), semantically and associatively related (*dog-cat*), only semantically associated (*bread-cereal*), or unrelated (*beard-tower*). They found that the purely semantically related pairs evoked a neural response in the brain that was the same as the response evoked by unrelated pairs of words. Associatively related pairs decreased the magnitude of the N400 effect,<sup>5</sup> whether the pairs also had a semantic relationship or not (see Figure 3.4). Recent behavioral and neuropsychological studies also suggest that people respond differently to association than they do to semantic relatedness. People respond to association more quickly than they respond to semantic relatedness (Perea and Rosa, 2002). Alzheimer's dementia patients also show priming for associatively related words, but not for solely semantically related pairs like *bread-cereal* (Glosser and Friedman, 1991; Glosser et al., 1998; Ober et al., 1995).

Concepts that co-occur more often in real life can become more strongly connected in the semantic network. As Perea and Rosa (2002, p. 189) explain, "the terms for things frequently connected in experience become themselves connected in the mental lexicon."

Patterns of connectivity between different words affect how easy it is to remember words. *Connectivity* reflects how many words are associated with a specific target word and how many connections are shared between that set of words (see Figure 3.5). Some



**Figure 3.4** ERP results for a priming experiment involving associatively related and semantically related pairs of words. The ERP waveforms in the box show that associated pairs decreased the magnitude of the N400 effect, but semantically related pairs did not. The response to semantically related pairs diverges from the response to the unrelated word pairs at a later point in time. *Source:* Rhodes and Donaldson (2008), John Wiley & Sons



**Figure 3.5** Connectivity for *dinner* and *dog*. *Source:* Nelson et al. (1993), American Psychological Association

words have few associates, and those associates have few connections between them. Those words have *low connectivity*. *High connectivity* words have more associates, and those associates have more connections between them. In Figure 3.5, *dog* is low in connectivity and *dinner* is high. High connectivity words are easier to remember than low connectivity words in both cued and free recall (Nelson et al., 1993; see also Breedin et al., 1998; Mirman and Magnuson, 2008). High connectivity words produce different patterns of brain activity in the temporal lobes than low connectivity words (Pexman et al., 2007; Wible et al., 2006). Thus, the structure of associations in memory affects the degree to which processing one word facilitates processing of a subsequent word, memory for individual words, and the brain's response to different words.

### *Associationist accounts of word meaning: Hyperspace analog to language and latent semantic analysis*

Whether “pure” automatic semantic priming exists or not, associative relations seem to play a powerful role in how people respond to words, which suggests that associative relations are encoded in the lexical representation of word form, meaning, or both. Some mathematical models of semantic memory place great emphasis on pure association (the extent to which words co-occur in the linguistic environment) and propose that association itself forms the basis upon which word meanings are built (Burgess and Lund, 1997; Camacho-Collados and Pilehvar, 2018; Landauer, 1999; Landauer and Dumais, 1997; Landauer et al., 1998; Lund and Burgess, 1996; see Günther et al., 2019 for a thorough review). Two prominent models of this type have been developed in the past decade or so, Burgess and Lund's *Hyperspace Analog to Language* (HAL)<sup>6</sup> and Landauer and Dumais's *Latent Semantic Analysis* (LSA). (You can explore how LSA works for yourself at <http://lsa.colorado.edu>.)

According to HAL and LSA, a word's meaning depends on the words that it appears with. If two words appear together more than they appear with other words, then the meanings of those two words are related. To determine whether two words are associated, HAL and LSA both depend on *corpora*,<sup>7</sup> which are large collections of utterances, which ideally reflect random, representative samples of the utterances that appear in the language as a whole. HAL's corpus included over 200 million words that were taken from USENET, an internet resource that has chat groups on a wide range of topics. HAL tracks 70,000 different words and uses its corpus data to determine how likely it is that each word will appear in the same utterance as each other word. For each word pair, HAL assigns a co-occurrence value based on how close the two words are, up to a distance of 10 words. Words that appear adjacent to one another get a score of 10. Words that are separated by one word get a score of 9, and so forth. At the end of this process, HAL has a 70,000 by 70,000 matrix that reflects word-to-word co-occurrence. A word's meaning is defined as the pattern of values in each of the cells in the matrix for each word. So each word has 140,000 numbers assigned to it, and the pattern of numbers, the *vector*, is the word's meaning.

LSA's original corpus included almost five million words that were taken from an encyclopedia. LSA divided its corpus into 30,000 episodes and assessed the number of times each one of 60,000 words appeared in each episode. LSA, like HAL, starts with a matrix. But unlike HAL, LSA assesses the relationship between a word and a number of contexts or episodes, rather than directly measuring co-occurrence between different words. In LSA, a word that appears many times in episodes 1 and 29,000 would get a

high number in those two cells. Once the cell values have been assigned, LSA subjects them to a form of factor analysis that captures commonalities in patterns of co-occurrences between words and episodes. Instead of 30,000 individual values being assigned to each word, factor analysis reduces the number of values to about 300. Similar to HAL, a word's meaning is represented in LSA as a pattern of values (a vector) across the 300 dimensions.

HAL and LSA use different methods to assess the degree to which words co-occur, but they share the idea that semantic representations incorporate a large number of dimensions (hundreds, in fact) and that word meanings can be described as vectors across those large numbers of dimensions.

HAL has been used successfully to model priming effects in lexical decision (Lund et al., 1995) as well as how people categorize words (Burgess and Lund, 1997). LSA has successfully modeled judgments of semantic similarity (saying whether two words are synonyms or not), aspects of children's vocabulary development (Landauer and Dumais, 1997), and judgments about the quality of text summaries (Kintsch et al., 2000; León et al., 2006). A similar high-dimensional model of word meanings has been used successfully to predict which brain regions will become most activated in response to a particular word (Mitchell et al., 2008).

One of the advantages of high-dimensional co-occurrence approaches to semantics is that they avoid some of the problems associated with the feature-based approach to word meaning. We can ask people to list features of objects, and we can use those lists of features to predict reaction times and similarity judgments. On this account, semantic similarity is a function of the number of overlapping semantic features—words with more features in common have more similar meanings. But there is (currently) no objective way to decide whether the mental representation of a word actually includes all and only the features that people list when we ask them to introspect about words, and feature-based representational theories can always be modified to include new features in the face of unexpected experimental results, which makes such accounts difficult to falsify (Buchanan et al., 2001). LSA and HAL get around the problem of subjectivity in feature descriptions by doing away with subjective feature descriptions altogether. Their methods of calculating semantic similarity are entirely objective and, hence, replicable and falsifiable.

### *The symbol-grounding problem*

Although HAL and LSA make good predictions for similarity judgments and some aspects of categorization, some language scientists are not comfortable with the idea that meaning depends entirely on word-to-word associations, whether based on simple co-occurrence or mathematically transformed co-occurrence. The chief among these objections also applies to semantic network theory. As explained by Art Glenberg and others (e.g. Glenberg and Robertson, 2000; Harnad, 1990; Lake and Murphy, 2021; Zwaan and Rapp, 2006), co-occurrence and association are not sufficient, by themselves, to describe word meanings, because associationist approaches like HAL and LSA merely describe mappings, albeit highly complex mappings, between symbols. Unless those symbols are *grounded* in some set of representations outside the symbol system, the symbols cannot be assigned any meaning.

There are different versions of this position. One of them is John Searle's *Chinese Room* argument (Searle, 1980). Searle asks you to imagine being an English speaker in

a small room with two slots in it, a rule book, and a stack of cards that have Chinese characters printed on them. You speak no Chinese and you do not know the meanings of any Chinese characters, but you do have a rule book that tells you what to do. When a Chinese character comes in one slot, you consult the rule book which tells you to pick some other characters out of your stack and push them out the other slot. If you have the correct rule book, you can respond perfectly appropriately, and your behavior would be entirely compatible with a native speaker of Chinese. So, if the characters coming in said "two plus two equals?" you could consult the rule book, and it would tell you to pick out the character that goes with "four." People from the outside would think that whoever is in the box understands every statement perfectly. But actually, you do not understand anything, you are just responding to symbols based on what the rule book says you should do. To you, those symbols have no meaning. The symbols could just as well have been, "IO," "F," "□O," and "‡." Until we *ground* those symbols in something other than more symbols, they have no semantic content, and hence no meaning.

Here's another way of looking at the *grounding* problem. Let's go back to our semantic network model for a moment. According to semantic network theory, the meaning of *goose* is based on a pattern of activation among a group of nodes associated with *goose* via links in the network. We understand what the *goose* node represents by seeing what nodes it is connected to, and what kinds of links connect the different nodes. But how do we understand the nodes that are connected to *goose*? We understand the meaning of those nodes by seeing what nodes they are connected to, and what kinds of links connect the different nodes. But how do we understand *those* nodes? By seeing what nodes they are connected to ... you get the idea.<sup>8</sup> The argument is equivalent to the *Chinese Room* case. Unless the symbols in the semantic network (or primitive feature network or high-dimensional-analog-to-language network) are connected to something other than abstract symbols, they can have no meanings. We might as well replace the labels in the semantic network with those in Figure 3.6. Or (as Art Glenberg put it many years ago) those in Figure 3.7.

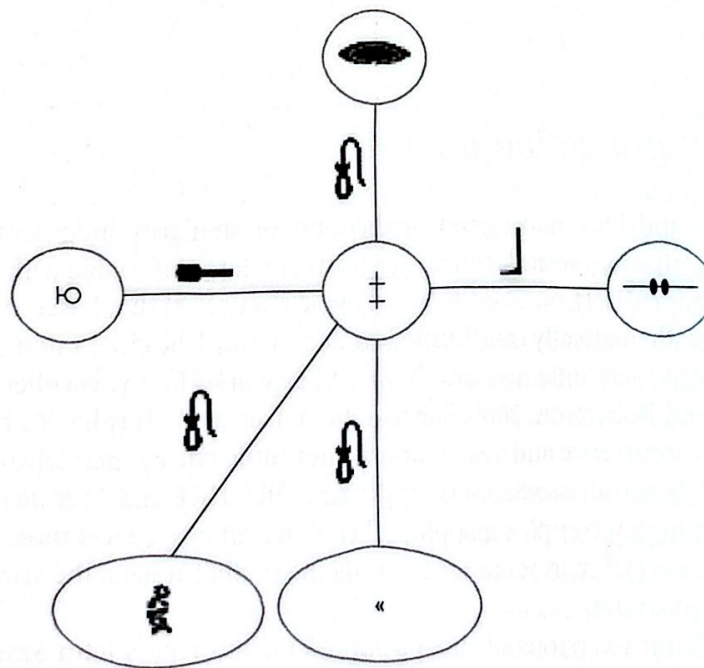


Figure 3.6 A hypothetical "semantic" network




Figure 3.7 Another hypothetical “semantic” network

### *Embodied semantics*

How can the grounding problem be solved? One answer is the *embodiment* or *embodied semantics* approach to meaning. Embodied semantics argues that abstract symbols or groups of symbols, like words, carry meaning because those symbols are tied to representations outside of the (traditionally defined) linguistic system. Specifically, words are tied to representations that we build using our perceptual apparatus (our five senses: vision, hearing, touch, taste, and smell). In that way, words do not just activate patterns of abstract symbols, words evoke perceptual experiences with real-world objects. When someone says *cat*, you do not just “think” “IO F ◻ O ‡,” you model the features of actual cats using the same apparatus that you use to perceive a real live, flesh-and-blood kitty. Glenberg and Robertson (2000) refer to this principle as the *indexical hypothesis*.

According to the indexical hypothesis, establishing a word’s meaning requires three processes. First, the word must be tied or *indexed* to actual objects in the world or analog representations of those objects in the mind (the *projected world* in Jackendoff’s, 1983, terminology; see also Speed and Majid, 2020; Zwaan et al., 2004). *Analog* representations are contrasted with *abstract* representations, in that analog representations carry some of the features of the actual object itself (Kosslyn, 1973; Stevens and Coupe, 1978). For example, a picture of a horse is an analog representation of a horse. The same information can be captured in an *abstract* way in a JPG file as a sequence of zeroes and ones. Glenberg refers to the analog mental representations of real-world objects as *perceptual symbols*, which implies that people have the ability to mentally manipulate these symbols as appropriate in the context that the utterance provides (Barsalou, 1999).

In the second step, people “use the indexed object or perceptual symbol to derive affordances” (Glenberg and Robertson, 2000, p. 384). The idea of *affordances* comes from the work of J.J. Gibson, a prominent theoretician and researcher in the area of perception. *Affordances* are determined by the interaction of our perceptual abilities and the physical characteristics of our bodies and the physical properties of objects in the world. For example, what makes a chair a chair is that the combination of our bodies’ physical

properties and the chair's physical properties allows us to use the chair for sitting—the chair *affords* sitting. (Chairs' physical properties provide other affordances as well. A chair can be used for self-defense in a bar fight. It can be used to raise the body and retrieve snacks from the high shelf, and so forth.)

The third process in creating meaning from utterances is to *mesh* or combine the affordances of the different indexed objects and characters in the utterance. When we interpret utterances, we index words to real-world objects by activating perceptual symbols, and the combinations of perceptual symbols determine what actions are available, how objects and actors might interact, and therefore what events are possible or likely.

Embodied semantics offers a potential solution to the symbol-grounding problem, but does that mean that the hypothesis is actually true? We can contrast predictions made by the indexical hypothesis with predictions generated from the high-dimension co-occurrence approaches to word meaning. To do so, Glenberg and Robertson (2000) constructed scenarios where critical objects were *not* good fits to the context based on co-occurrence metrics derived from HAL and LSA, but where the described actions *were* consistent with affordances derived from the situation described in the scenario. For example, subjects read this context sentence (Glenberg and Robertson, 2000, p. 385): *Marissa forgot to bring her pillow on her camping trip*. Subjects would then read one of two continuations. One of the continuations included a word that matched the affordances generated by the situation; and the other continuation did not. The afforded continuation was, *As a substitute for her pillow, she filled up an old sweater with leaves*. The non-afforded continuation was, *she filled up an old sweater with water*. *Leaves* matches the affordances of the described situation (because you really could fill up an old sweater with leaves), but *water* does not. Critically, the LSA association values for the two test words, *leaves* and *water*, in the context are the same. According to LSA, both continuations are equally meaningful and good. When people judged the plausibility of the two continuations, though, they rated the afforded continuation as being much better than the non-afforded continuation. In this case, the indexical hypothesis, but not the high-dimensional symbol association approach, accurately predicted how people would judge the meaning of the sentences.

Since the original work on the indexical hypothesis was published, a number of other studies have pointed toward a relationship between the linguistic-semantic system and perceptual and motor systems that have been traditionally viewed as outside the language system (see Davis and Yee, 2021). Specifically, the perceptual and motor systems may provide some of the machinery that creates meaning within the linguistic system. But apart from off-line judgment tasks like Glenberg and Robertson's (2000), what evidence suggests that the semantic system relies on perceptual and motor systems? In fact, there is growing evidence that they are.

In one study, participants made plausibility judgments (i.e. "Yes or no, does this statement make sense?") after reading statements like *He opened the drawer*, and *He closed the drawer* (Glenberg and Kaschak, 2002; but see Morey et al., 2021 for a large-scale failure to replicate). The experimenters manipulated whether subjects responded by moving their hand away from their body to press a key or moved their hand toward their own body. In some conditions, the sentences implied motion toward the body (opening a drawer means you pull the drawer towards yourself) or away from the body (closing a drawer indicates motion away from yourself). The subjects' answers could be made with a body movement that was either the same as the motion implied by the sentence (move your hand away from yourself to answer "yes"; the sentence says *He closed the drawer*), or the subjects could answer with a body movement that was opposite to that indicated in the sentence (move your hand away to answer "yes" to *He opened the drawer*). In this experiment, subjects' responses were faster when the motion undertaken to answer the question matched the motion indicated by the sentence. Why did this happen?

According to the embodied semantics position, people understand the meaning of expressions like *open/close the drawer* by indexing the words to perceptual symbols (mental models of the objects), and then mentally simulating the action indicated by the sentence. To mentally simulate the action in the sentence, you use the same motor system that you use to move your actual body. Hence, your physical response to the question uses the same resource that you use to figure out the meaning of the sentence—the motor system. If word meanings were based on arbitrary and abstract networks of symbols, there is no reason why language should have any effect on your body movements. If meaning is governed entirely by abstract symbol systems, you should respond just as quickly no matter what direction you need to move your hand.

Also, responses to a word speed up when people's hands are shaped like they would be if they were actually using the named object, such as a pen or a knife, suggesting that action in the motor system can facilitate response to a word (Klatzky et al., 1989; see also Lieberman, 2000; Setola and Reilly, 2005). Additional studies provide further evidence that individual word meanings and motor responses interact (Tucker and Ellis, 2001, 2004; see also Buccino et al., 2005; Zwaan and Taylor, 2006; but see Miller et al., 2018).

One kind of study capitalizes on the fact that we interact with some objects using a *precision grip*—like pens, silverware, and buttons. We interact with other objects—like hammers and shovels—using a *power grip*. Subjects read words presented one by one on a computer screen and judged as quickly as possible whether the object was natural or man-made (a kind of *semantic categorization* task). Half of the participants made their response using a power grip, and half made their response with a precision grip. Subjects in the power-grip response condition responded faster to words describing power-grip objects and slower to words describing precision grip. Subjects in the precision-grip response condition showed the opposite pattern.

Other experiments showed that the word–motor connection goes in the other direction as well—word processing has an effect on the motor system. When people make a motor movement at the same time as they are reading an action-related word, motor movements are slowed down. But if people read the same words *before* they begin moving, movements are speeded up (Boulenger et al., 2006). These results can be explained if we assume that part of a word's meaning includes a mental simulation of the object, and that these mental simulations involve modeling how you typically move as you interact with the object.

The word–action compatibility effects just described suggest a close relationship between a system of meanings and the motor system that we use to move ourselves around and interact with objects in the world. The results of some ERP, *transcranial magnetic stimulation* (TMS), and single-cell recording experiments support this view as well. In ERP studies that seek to locate the source of electrical activity in the brain, the strongest response to action–word stimuli occurs at parts of the scalp that are directly over the motor strip—the part of the brain that is responsible for planning body movements (Aflalo et al., 2020; Dreyer and Pulvermüller, 2018; Hauk et al., 2004; Pulvermüller et al., 2001; Pulvermüller et al., 1999; Pulvermüller, Shtyrov et al., 2005; Vitale et al., 2021).<sup>9</sup>

Using a technique similar to ERP, *motor evoked potentials* (MEPs) can be measured based on the activity at *neuromuscular junctions*—the places where efferent nerves connect with skeletal muscle tissue. Neuromuscular junctions are responsible for activating muscles in response to signals from the brain. When people listen to sentences that describe hand-related actions (like sewing), MEP activity measured above hand muscles decreases, but foot muscles are unaffected. When people listen to sentences about foot-related actions, the opposite pattern occurs (Buccino et al., 2005). Finally, when Parkinson's dementia patients were tested, they showed reduced priming for action-related

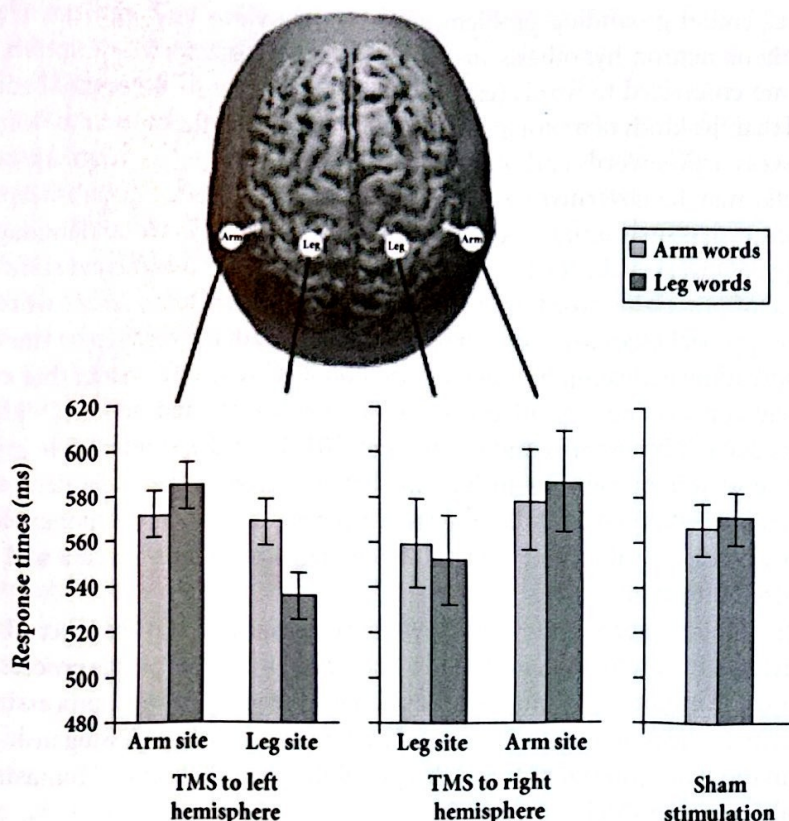
words but normal priming for other kinds of words when they were *not* taking medication that improves motor function. When the patients were on medication that boosts motor function, the differences in amount of priming between action words and non-action words disappeared (Boulenger et al., 2008).

In TMS experiments, a powerful magnetic field is generated very close to the scalp, which can induce electrical activity in populations of neurons directly beneath the TMS device. TMS-induced neural activity can facilitate the mental processes that are involved in information processing (perhaps by synchronizing neural activity in a population of neurons involved in the task), and so TMS can be used to assess the extent to which different parts of the brain participate in different aspects of language processing.

In one TMS study, people made lexical decisions in response to words that were related either to arm movements or leg movements (Pulvermüller et al., 2005). The words were presented in writing on a computer screen and 100 ms after the words appeared on the computer screen, TMS was applied either over parts of the motor cortex involved in hand movements or parts involved in leg movements. In a control condition, *sham* (fake) TMS was applied. As illustrated in Figure 3.8, TMS speeded lexical decisions to words related to the part of the body controlled by the part of motor cortex that was stimulated. When arm areas were stimulated, lexical decisions to arm-related words were made more quickly. Semantic processing of words is thought to be left lateralized, and in this study TMS affected responses to words when it was applied over the left hemisphere, but not when it was applied over the right. Sham stimulation also had no effect, which rules out demand characteristics as the source of the reaction-time effects. Other TMS studies produced comparable results. For example, Oliveri and colleagues used TMS to show that motor cortex responds more strongly to action-related nouns and verbs (*the axe, to bite*) than to nouns and verbs that do not have specific associated actions (*the cloud, to belong*) (Oliveri et al., 2004).

fMRI investigations of the relationship between word processing and the motor system have capitalized on the fact that the motor system is organized such that different parts of the brain are responsible for controlling different parts of the body. If there is a close connection between the linguistic-semantic system and the motor system, and if the semantic system “borrows” parts of the motor system to instantiate the meanings of particular words, then different kinds of words should produce different patterns of activation in the motor system.

When fMRI was used to determine where activity in the brain occurred in response to words that refer to movements of the face (*smile*), arm (*throw*), or leg (*walk*), increased activity was observed in “classical” language areas, like *Wernicke’s* area (at the junction of the parietal, occipital, and temporal lobes in the left hemisphere) and *Broca’s* area (in the left frontal lobe, just in front of the motor strip; Hauk et al., 2004), but increased activity was observed in other brain areas as well. The neural response to words referring to body movements was compared to brain activity that occurred when subjects actually moved the corresponding body part (see Plate 1). The striking result here is that words related to actions led to increased neural activity in the same parts of the brain that became active when subjects actually moved the corresponding body part. This result is consistent with the embodied semantics view that word processing involves the activation of perceptual-motor representations. Listening to sentences describing face, arm, and leg actions produces a similar pattern of activation, with areas of motor and premotor cortex activated to different degrees by face-, arm-, and leg-related action sentences (Tettamanti et al., 2005). Reading sentences has comparable effects. The parts of the brain that become active when a person views an action also become active when the same person reads a sentence describing the corresponding action (Aziz-Zadeh et al., 2006).



**Figure 3.8** TMS and lexical decisions. The top picture shows where TMS was applied in the left and right hemispheres. Response times on the lexical decision task appear below the brain. Left-hemisphere stimulation affected lexical decision latencies, but right-hemisphere stimulation did not. In the left hemisphere, arm words were responded to more quickly following TMS over the part of the motor cortex that controls arm movements. A similar effect was observed for leg-related words after leg-area stimulation. Sham TMS had no effect. *Source:* Pulvermüller, F. et al. (2005), John Wiley & Sons

Some scientists believe that *mirror neurons* provide the neural basis for mental simulations proposed by embodied semantics in addition to the motor system. Mirror neurons become active when monkeys engage in an action—like grasping a cup—and they also become active when monkeys watch someone else engage in the same action (Gallese and Lakoff, 2005; Rizzolatti and Arbib, 1998; Rizzolatti and Craighero, 2004). Further, the part of the monkey’s brain that is analogous to Broca’s area, a part of the frontal lobe classically associated with speech, contains mirror neurons (Buccino et al., 2005). The idea is that the linguistic–semantic system in humans also drives mirror neurons and uses them to represent the meanings of words that describe objects and actions. On this account, perceiving the word *hammer* triggers a response in the mirror neuron system that closely resembles the pattern of neural response that happens when we use a hammer ourselves or watch someone else use a hammer. Research in word processing shows that merely observing a hand shape has similar effects on word processing as actually making the hand shape, in particular on identifying what category a word belongs to (as in Klatzky et al., 1989; Vainio et al., 2008). These results are consistent with the mirror neuron hypothesis because observation appears to have similar effects on the interaction between the motor and language systems as real action does.

Although the embodied semantics approach does a good job explaining why and how words affect motor regions of the brain and vice versa, and it goes a long way toward

solving the symbol-grounding problem, not everyone views embodiment in general, and the mirror neuron hypothesis in particular, as a satisfactory description of how meanings are connected to words (e.g. Ostarek and Huettig, 2019). Some theorists are concerned that the kinds of motor and perceptual simulations that appear to occur when people process action words and phrases (as indicated by response–language compatibility effects) may be governed by a separate system than the language interpretation system, and/or that such simulations may be an optional component of language interpretation (e.g. Oliveri et al., 2004). Other researchers propose that mental simulation is a byproduct of processing words and is not strictly necessary to represent word meanings. Other approaches view activation of the motor system by words to be the result of a kind of spreading activation between a “disembodied” semantic system that connects to separately functioning cognitive systems for perception and action (Mahon and Caramazza, 2008; Vannuscorps and Caramazza, 2019). Finally, some people argue that the existence of mirror neurons in humans has not been conclusively demonstrated (M. A. Gernsbacher, 2009), and that the mirror neuron hypothesis cannot explain why damage to Broca’s area does not lead to comprehension deficits (Corina and Knapp, 2006; see also Lotto et al., 2008).

If motor simulation of actions is an inevitable consequence of word processing, then neural activity in the motor system should be observed whenever people process action-related words. If motor simulation is an optional byproduct of word processing, then neural activity in the motor system may occur after some word-processing tasks but not others. Tomasino and colleagues tested this possibility in a TMS study (Tomasino et al., 2008; see also Montero-Melis et al., 2022).

Tomasino and colleagues zapped their subjects with TMS pulses while they were processing action words. They manipulated the lag between presenting the word and applying a TMS pulse as well as the experimental task. When participants engaged in an explicit visual imagery task (subjects imagined themselves performing the action denoted by a target word and said whether the action required wrist rotation), TMS facilitated the response, and then only when the TMS pulse was delivered about 90 ms after the target word. Other tasks with the same target words, silent reading, and frequency judgment were not affected by the TMS pulses. Thus, previous positive findings in TMS studies may reflect an optional element of visual imagery, rather than reflecting necessary outcomes of word processing. Additionally, different types of language may evoke motor representations to different degrees. For example, figurative language (such as metaphors) may not evoke spatial models in the same way that literal language does, which calls into question the universality of perceptual simulation in word processing (Bergen et al., 2007).

Although some neuropsychological data support the integration of linguistic and (traditionally defined) nonlinguistic systems for action comprehension (e.g. Saygin et al., 2004, see Plate 2), neuropsychological data from patients with brain damage can also be used to argue that the semantics of action words does not depend on perceptual-motor representations. First, lesions in motor cortex are not always followed by problems recognizing and understanding action words (Argiris et al., 2020; De Renzi and Di Pellegrino, 1995; Saygin et al., 2004). When Saygin and colleagues measured the relationship between lesion location and degree of impairment on different tasks, they found some regions of the brain that caused impairment for reading of action-related words, but that did not associate with impairment for perceiving those actions. Other regions correlated with impairment of action perception, but not reading about actions. This suggests a separation between the linguistic–semantic system and the visual-perceptual system. Saygin and colleagues therefore suggested (p. 1799) that, “There was no overall correlation between patients’ deficits in the two domains [visual perception

and reading], suggesting that the deficits observed in actions and comprehension of actions through processes.”<sup>10</sup> Negri and colleagues (2007) also show that actions and the knowledge necessary to recognize actions are dissociable. Some people can recognize actions that they are unable to perform. Some people with damage in the motor area. This calls into question whether the ability to mentally simulate actions is a necessary component of recognizing actions (see also Mahon and Caramazza, 2008).

To summarize, the semantic network model is a promising approach to action semantics. Connections between words and the processes that they activate can help explain why different patterns of priming occur in a variety of experimental tasks. HAL and LSA projections in the semantic network capture the essence of the semantic network, but does not offer an answer to the symbol-grounding problem. Whether perceptual simulation offer a potential answer to the symbol-grounding problem there is a growing body of experimental evidence on the role of word processing and parts of the brain that are involved in these processes. Whether those connections are necessary for action recognition is an open question.

## Lexical Access

*Most models of lexical access do not actually deal with the problem of how words are accessed.*  
GARETH GASKELL AND WILLIAM MARSLÉN-WILSON

*Lexical access* involves the set of mental representations that we use to identify specific words as we are listening or reading. It is the activation of semantic information, but models of lexical access usually focus on the activation of word *form* information (what the word sounds like or what they look like), with the activation of semantic information as a consequence of the activation of form. The activation of semantic information in spoken language processing is automatic and seems to require no conscious effort. For many people, it is really nothing there to explain. For many people, it is a conscious effort. This apparent ease and automaticity of lexical access involves complex mental operations and, as a result, the debate continues among language scientists about how lexical access is organized, what exact mental mechanisms are involved in lexical access, what exact mental mechanisms are involved in lexical access, what exact mental mechanisms are involved in lexical access.

Let’s start with the fact that people are able to identify words quickly. In seminal work in this area, William Marslen-Wilson and colleagues (1980) used a shadowing task to estimate how much time people need to identify words. In a shadowing task, subjects listen to recorded speech and repeat the words as quickly as possible the words that they hear. Spoken words are presented for a second (at an average speaking rate of 158 words per second). This means that they are following along approximately 100 ms behind the speaker.

When Marslen-Wilson analyzed the kinds of errors that people made in the shadowing task, errors were not random and did not consist of

and reading], suggesting that the deficits observed in the comprehension of pantomimed actions and comprehension of actions through reading are not tightly coupled processes.”<sup>10</sup> Negri and colleagues (2007) also showed that knowledge of how to produce actions and the knowledge necessary to recognize actions do not always go together. Some people can recognize actions that they are not able to produce because of brain damage in the motor area. This calls into question the idea that using the motor cortex to mentally simulate actions is a necessary component of recognizing and understanding actions (see also Mahon and Caramazza, 2005).

To summarize, the semantic network model is still a standard theory of lexical semantics. Connections between words and the process of automatic spreading activation help explain why different patterns of priming occur for different kinds of words across a variety of experimental tasks. HAL and LSA propose that the structure of the associations in the semantic network capture the essence of word meanings, but that position does not offer an answer to the symbol-grounding problem. Embodied semantics and perceptual simulation offer a potential answer to the symbol-grounding problem, and there is a growing body of experimental evidence that indicates a connection between word processing and parts of the brain that are responsible for perceptual and motor processes. Whether those connections are necessary for word understanding remains an open question.

## Lexical Access

*Most models of lexical access do not actually deal with activation of meaning.*

GARETH GASKELL AND WILLIAM MARSLÉN-WILSON

*Lexical access* involves the set of mental representations and processes that allow us to identify specific words as we are listening or reading. Recognizing words leads to the activation of semantic information, but models of lexical access typically deal specifically with the activation of word *form* information (stored representations of how words sound or what they look like), with the activation of semantic information being treated as a consequence of the activation of form. The recognition of familiar words during spoken language processing is automatic and seemingly effortless. It may seem like there is really nothing there to explain. For many people, but certainly not all, reading seems similarly effortless. This apparent ease and automaticity obscures the fact that lexical access involves complex mental operations and, despite its apparent ease, considerable debate continues among language scientists about which exact properties of words are involved in lexical access, what exact mental mechanisms take part, and how the entire process is organized.

Let's start with the fact that people are able to identify spoken words really amazingly quickly. In seminal work in this area, William Marslen-Wilson (1973) employed a *shadowing task* to estimate how much time people needed to identify spoken words. In the shadowing task, subjects listen to recorded speech and they try to repeat (or shadow) as quickly as possible the words that they hear. Spoken language delivers about 5 syllables per second (at an average speaking rate of 158 words per minute). Some people, *fast shadowers*, are able to repeat a stream of words at a lag of as little as 250 ms (a quarter of a second). This means that they are following along approximately one syllable behind the input.

When Marslen-Wilson analyzed the kinds of errors that people made, he found that errors were not random and did not consist of mere pronunciation difficulties. Instead,


when people made errors, the incorrect words that they produced were fully compatible with the semantic and syntactic content of the preceding context. Out of 132 errors where subjects replaced or added words, only three violated syntactic constraints on acceptable continuations. This means that fast shadowers were able to perform lexical access very fast indeed, and that they computed higher order aspects of the speech stream—minimally, its syntactic form—within a few hundred milliseconds of the word's onset. Findings like these show that speech processing and lexical access from spoken words are both highly *incremental*—the speech stream is segmented into words, and higher order relationships between words are represented before major clause or sentence boundaries are encountered. To explain how people process spoken words, the very least we need is a system that can identify individual words very quickly.

Additional evidence for very fast lexical access comes from *word monitoring* and *gating* tasks (Grosjean, 1980; Marslen-Wilson, 1987; Marslen-Wilson and Tyler, 1980). Word monitoring involves listening to speech and responding as quickly as possible when a specific target word appears in the input. The *gating* task involves listening to short snippets of the beginnings (*onsets*) of words. The subject's task is to say what word is present in the stimulus. The length of the snippet is increased by small increments (25 or 50 ms) until the subject can correctly say what word the snippet belongs to. The length of the snippet serves as an estimate of how much *bottom-up* information (auditory stimulation) the subject requires to identify the word. These different tasks all provide roughly the same estimate of the amount of input it takes for people to identify spoken words. For one- and two-syllable words in the context of a spoken sentence, the average is about 200 ms worth of input (a fifth of a second; Marslen-Wilson, 1973, 1985; Marslen-Wilson and Tyler, 1975; Seidenberg and Tanenhaus, 1979); it takes about another 100 ms of input before people can recognize isolated words.

Word forms can be divided up and analyzed according to their components. Spoken words can be divided into phonemes, which can be further divided into phonetic features. We can also view words as being made up of syllables, which in turn are composed of sets of phonemes. Words can also be thought of as being made up of organized *sub-lexical* ("below the level of the word") units called *morphemes*. Different theories of lexical access make different claims about which of these units affect the process of recognizing specific words from spoken input. Some theories propose that phonetic features, but not phonemes, play a role. Some theories propose that phonetic features, phonemes, and word-level representations all play a role. Some theories propose that word meanings themselves play a role in lexical access. To organize the discussion, this section starts with first-generation accounts including *logogen*<sup>11</sup> and the *frequency ordered serial bin-search* models (Clarke and Morton, 1983; Jackson and Morton, 1984; Morton, 1969; Taft and Forster, 1975). Then it turns to second generation accounts, such as the original version of the *COHORT* and *TRACE* models.

Finally, it discusses third-generation accounts, such as the *distributed cohort model* and the *Simple Recurrent Network* approach.

All of these accounts have a common goal: They try to explain how people take inputs from the auditory or the visual system and match those inputs to representations of word form in long-term memory. To explain how that is done, a theory of lexical access has to say how the mind organizes the input—what characteristics or features it perceives in the input—and how it connects those characteristics to word form representations. As a starting point, consider this *default model* of lexical access. Words are made up of parts. Some of those parts (e.g. phonetic features) are more basic than others (e.g. syllables). We could have a model of lexical access that says: Take a segment of speech, start by identifying the most basic units (e.g. phonetic features), combine those features to find more complex units (e.g. phonemes), combine those features to find even more complex



**Figure 3.9** A hypothetical bottom-up model of lexical access (for simplicity, only some of the possible connections are illustrated). Information flows in the direction indicated by the arrows

units (e.g. syllables) and then use those units to find stored words that have matching forms. This is called a *bottom-up* processing system, because information flow in the system starts with more basic units, which are conceived of as being at the bottom of a hierarchy like the one in Figure 3.9, and proceeds upwards through more and more complex units. Information could flow from higher level representations to lower level ones, which is called *top-down* processing. Models of lexical access differ from one another in terms of the kinds of representations that they believe participate in lexical access as well as the way information flows throughout the system.

### First-generation models

John Morton's *logogen model* is a bottom-up driven system that takes spoken or visual input and uses it to activate previously stored word form representations (Morton, 1969). The heart of the logogen model was a set of processing units that would receive input from either spoken or written modalities and would fire when their excitatory inputs exceeded some criterion level or *threshold*. As Morton notes (p. 165): "The logogen is a device which accepts information from the sensory analysis mechanisms concerning the properties of linguistic stimuli and from context producing mechanisms. When the logogen has accumulated more than a certain amount of information, a response (in the present case the response of a single word) is made available."

Max Coltheart and his colleagues describe the logogen system in this way (Coltheart et al., 2001, p. 209, modern day classifiers might implement the logogen idea), "*Logogens* are evidence-collecting devices with thresholds. Evidence is collected from visual or auditory input, and when the amount of evidence collected by a word's logogen exceeds that logogen's threshold, information about that word in the cognitive system (e.g. its meaning) is accessed."

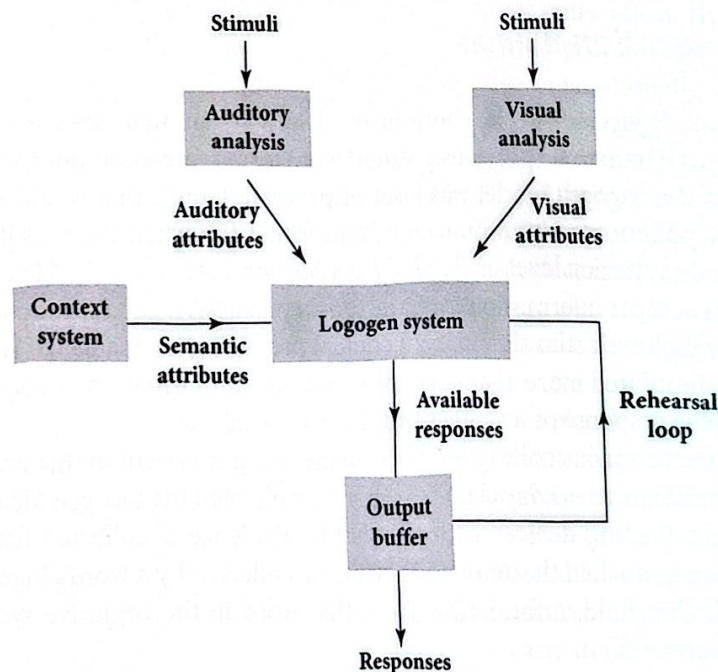
Each word in a person's vocabulary is represented by a logogen; and words are recognized when the activation levels of their corresponding logogen exceed some threshold. What has to happen to make a logogen exceed its activation threshold? In Morton's

system, logogens receive inputs from either spoken words, written words, or preceding context (which would activate logogens based on semantic attributes). Normally, input would come from either the auditory or visual systems, and not both at once, but both auditory and visual evidence could be present simultaneously (as in reading subtitles on Netflix). The semantic input mechanism allows for context to influence the amount of time it takes to recognize a word. Context words that are semantically related to an individual logogen will raise the activation of the logogen before the listener gets direct perceptual evidence that the corresponding word is actually present in the input.

The logogen system operates on these three kinds of inputs and, when individual logogens become activated at a level above their thresholds, they send signals to an output buffer (see Figure 3.10). Unless new input continues to activate the logogen, a decay function returns its activation to baseline levels within about one second. Once a logogen has been triggered or activated, its threshold for activation is temporarily lowered. As a result, less evidence is needed in the acoustic and visual input channels to reactivate the logogen. This mechanism can account for repetition priming effects—it is easier to recognize a word the second time you see it than the first because the activation threshold is lower the second time around.

The logogen model makes two key assumptions. First, it assumes that information flow is strictly bottom-up. Auditory and visual processing units affect the activation of logogens, but logogens do not affect the activation levels of the auditory and visual processing units that feed into the logogen. Second, it assumes that there are no direct connections between and among the logogens themselves. As a result, the activation level of one logogen does not affect the activation of any of the other logogens.

The logogen model is an important one in psycholinguistics, because it was one of the first attempts to mathematically model (and thereby explain) how people respond to words. The model was successful in a number of ways. First, it had been known for a long time that word frequency affects a variety of behaviors. Words that occur frequently in a language take less time to recognize than words that appear less frequently. Why



**Figure 3.10** A schematic of the information flow in John Morton's *logogen model*.  
Source: Morton, J. (1969), American Psychological Association

should this be the case? Morton suggested that repeated exposure to high-frequency words lowers the threshold for activation in the logogens that represent those high-frequency words. So less external evidence (“bottom-up” input) is required before you can recognize a high frequency word, and therefore you respond faster to high-frequency words than low-frequency words. This may also help explain why high-frequency words tend to be shorter than lower frequency words (as per *Zipf’s Law*; Zipf, 1949). Shorter words pack less phonological and/or orthographic information than longer words, but this does not make them harder to recognize and process, because more frequent exposure lowers their activation thresholds. The model also helps explain why high-frequency words are easier to recognize than low-frequency words when they have been degraded by noise. Noise in the signal decreases the quality of the bottom-up input, but high-frequency words don’t need as much bottom-up input, so they are recognized even in noisy environments.

### MORPHOLOGY AND LEXICAL ACCESS

The logogen model was the first one to mathematically model the mental processes involved in lexical access (word form recognition) but was followed up shortly by other models. One of the most prominent subsequent models was Ken Forster and Marcus Taft’s *frequency ordered bin search* (FOBS) model (Forster, 1989; Forster and Bednall, 1976; Taft and Forster, 1975). Like logogen, FOBS proposed that word form representations were activated by bottom-up input from the auditory system. According to Taft and Forster’s model, people perform lexical access by using auditory (or visual) cues to search their long-term memories for a matching stimulus. This search process is organized so that people do not need to search the entire lexicon every time they need to look up a word. Instead, lexical (word form) representations are organized into *bins*. The bins are organized according to word frequency. High-frequency words are at the “front” of the bin and are searched first; lower frequency words are stored toward the “back” of the bin and are searched later. When you encounter an auditory stimulus, that opens up a bin (kind of like opening up a file drawer), and you search through the bin looking for an entry that matches the stimulus, starting with the most frequent item in the bin, then the next most frequent, and so on until you have searched the entire bin. The search process ends when you find an item in the bin that matches the stimulus. This kind of search is called *self-terminating* (the process stops itself when it succeeds), so you don’t keep searching the bins for an additional match after you have found one good candidate. One last important characteristic of the model is that words are organized in the bins according to shared *roots*. To define what a shared root is, and why it might be important, we need to discuss a bit of *morphology*.

The FOBS account proposes that *morphemes* are an important level of representation in lexical access, so we need to know what a morpheme is. Morphemes are defined as the smallest unit of language that can be assigned an independent meaning. Words are made up of one or more morphemes. The basic morpheme in a word is its *root*, or *root morpheme* (sometimes called a *stem*). *Board* is a *monomorphemic* word, because it cannot be decomposed into smaller units of meaning. So the root morpheme for *board* is the same as the word itself. *Chalkboard* is a *polymorphemic* word, because it can be decomposed into smaller units. Specifically, it contains the morphemes *chalk* and *board*, each of which has a meaning of its own, and each of which contributes to the meaning of the word as a whole. Which of the morphemes in *chalkboard* is the root? While some linguistic theories would argue that *board* is the root (because a *chalkboard* is a kind of a *board*, not a kind of *chalk*), the FOBS account proposes that *chalk* is the root, because speech processing gives priority to information coming first, and in speech, we hear the morpheme *chalk* before we hear the morpheme *board*.

The category *morpheme* can itself be divided into subcategories. Polymorphemic words are made up of a *root* and one or more *affixes*. (Compound words are special because they are made up by combining two or more *root* morphemes.) *Affixes* can be *prefixes* that come before the root, *suffixes* that come after the root, or *infixes* that divide a root into two parts, one of which comes before the infix, and one of which comes after. Standard American English does not have any infixes;<sup>12</sup> Arabic has many. Affixes make up the class of *bound* morphemes (as opposed to *free* morphemes), because they cannot appear by themselves (whereas free morphemes can). Affixes come in different types, as well. *Inflectional morphemes* change the flavor of a word's meaning; and *derivational morphemes* change the syntactic category that a word belongs to. We can change the flavor of the word *cat* without changing its core meaning or its syntactic category by adding the bound-morpheme, *-s*. We can change the tense of a verb by adding inflectional morphemes like *-ed* or *-ing* (e.g. *bake, baked, baking*). If we want to change the *category* of a word, we can add derivational morphemes like *-ly* or *-tion*. We can take a verb like *confuse* and change it to a noun with the *-tion* derivational morpheme—*confuse* becomes *confusion*. We can change the verb to an adjective with the *-ing* derivational morpheme—*confuse* becomes *confusing*. We can stack morphemes end to end to change from a noun to an adjective and back to a noun again—*truth* (n.) becomes *truthy* (adj.) becomes *truthiness* (n.). The morphological system in English is one of the properties that contributes to the productivity or generativity of the language. We can combine old morphemes in new ways to come up with new meanings.<sup>13</sup>

What do morphemes have to do with lexical access? It depends on how you think word representations are organized, and what you think happens when people encounter a polymorphemic word. The FOBS model says that lexical representations are organized into bins, and each bin is built around a root. All of the variants of *dog* are listed under a bin, and *dog* is the base entry.<sup>14</sup> So, *dog, dogs, dogged, dogpile, and dog-tired* are all represented in the same bin. Any time you encounter the root *dog*, you search through the *dog* bin looking for a matching entry.

Alert readers will have noted that there are many versions of *dog* that are not identical to the label on the bin (*dog*). What happens when the stimulus does not match the label on the bin? According to the FOBS model, the incoming stimulus has to be analyzed according to its root, because the root is what gets the listener access to the correct bin. Whenever a listener encounters a polymorphemic word (*dogs, dogpile, dogaphobia*), the first thing the listener needs to do is figure out what the root is. Therefore, the first step in lexical access is *morphological decomposition*—the incoming stimulus needs to be broken down into parts that correspond to individual morphemes before the root can be identified. A word like *dogs* is analyzed as being made up of the root morpheme *dog* and the plural inflectional suffix *-s*.

What evidence suggests that lexical access involves morphological decomposition? Such evidence comes in various forms. First, as noted previously, people respond more quickly to frequent words than infrequent words. But it's actually a bit more complicated than that, because we can measure frequency in different ways. We could assign frequency estimates to an entire word, regardless of how many morphemes it contains. We could look at a corpus and count up every time the word *dogs* appears in exactly that form. We could count up the number of times that *cats* appears in precisely that form. In that case we would be measuring *surface frequency*. But the words *dogs* and *cats* are both related to other words that share the same root morpheme. We could decide to ignore minor differences in surface form and instead concentrate on how often the family of related words appears. If so, we would treat *dog, dogs, dog-tired, and dogpile* as being a single large class, and we would count up the number of times any member of the class appears in the corpus. In that case, we would be measuring *root frequency*—how often

the shared word root appears in the language. Those two ways of counting frequency can produce very different estimates. For example, perhaps the exact word *dog* appears very often, but *dogpile* appears very infrequently. If we base our frequency estimate on surface frequency, *dogpile* is very infrequent. But if we use root frequency instead, *dogpile* is very frequent, because it is in the class of words that share the root *dog*, which appears fairly often.

If we use these different frequency estimates (surface frequency and root frequency) to predict how long it will take people to respond on a reaction time task, root frequency makes better predictions than surface frequency does. People respond quickly to words with low surface frequency if their root frequency is high (Bradley, 1979; Taft, 1979, 1994). This outcome is predicted by an account like FOBS that says that word forms are accessed via their roots, and not by models like logogen where each individual word form has a separate entry in the mental lexicon.

Further evidence for the morphological decomposition hypothesis comes from priming studies involving words with real and *pseudo-affixes*. Many polymorphemic words are created when derivational affixes are added to a root. So, we can take the verb *grow* and turn it into a noun by adding the derivational suffix *-er*. A *grower* is someone who *grows* things. A lot of words end in *-er* and have a similar syllabic structure to *grower*, but are not real polymorphemic words. For example, *sister* looks a bit like *grower*. They both end in *-er* and they both have a single syllable that precedes *-er*. According to the FOBS model, we have to get rid of the affixes before we can identify the root. Anything that looks or sounds like it has a suffix is going to be treated like it really does have a suffix, even when it doesn't. Even though *sister* is a monomorphemic word, the lexical access process breaks it down into a *pseudo-* (fake) *root*, *sist*, and a *pseudo-suffix*, *-er*.

After the *affix stripping* process has had a turn at breaking down *sister* into a root and a suffix, the lexical access system will try to find a bin that matches the pseudo-root *sist*. This process will fail, because there is no root morpheme in English that matches the input *sist*. In that case, the lexical access system will have to re-search the lexicon using the entire word *sister*. This extra process should take extra time, therefore the affix stripping hypothesis predicts that *pseudo-suffixed* words (like *sister*) should take longer to process than words that have a real suffix (like *grower*). This prediction has been confirmed in a number of reaction time studies—people do have a harder time recognizing pseudosuffixed words than words with real suffixes (Lima, 1987; Smith and Sterling, 1982; Taft, 1981). People also have more trouble rejecting pseudowords that are made up of a prefix (e.g. *de*) and a real root morpheme (e.g. *juvenate*) than a comparable pseudoword that contains a prefix and a non-root (e.g. *pertoire*). This suggests that morphological decomposition successfully accesses a bin in the *dejuvenate* case, and people are able to rule out *dejuvenate* as a real word only after the entire bin has been fully searched (Taft and Forster, 1975). Morphological structure may also play a role in word learning. When people are exposed to novel words that are made up of real morphemes, such as *genvive* (related to the morpheme *vive*, as in *revive*) they rate that stimulus as being a better English word and they recognize it better than an equally complex stimulus that does not incorporate a familiar root—such as *gencule* (Dorfman, 1994, 1999).

FOBS is also consistent with experiments showing that words that are related via a shared morpheme prime one another (Drews and Zwitserlood, 1995; Emmorey, 1989; Stanners et al., 1979; see also Deutsch et al., 2003).<sup>15</sup> In these priming experiments, the targets are root words, like *honest*, and the prime words are either identical to the target (*honest*) or a prefixed version of the target (e.g. *dishonest*). To control for possible effects at other levels, like letter overlap, in another condition the target would be a word like *son* and the prime would be a word that had many of the same letters, like *arson* (but the

two words do *not* share a root morpheme, because *arson* is not a kind of *son*). In these experiments, equivalent priming occurred when the prime was either identical to the target (e.g. *honest-honest*), or contained the target as a root (as in the *dishonest-honest* case). No priming was observed for words that only had overlapping letters (the *arson-son* case). These effects are compatible with FOBS, because prefixed words like *dishonest* are accessed via their roots. So, processing the word *dishonest* entails activating the representation of the root morpheme *honest*. If *honest* is presented right after *dishonest*, its lexical entry should be more activated than normal, which speeds up the response. Similar effects occur for words with suffixes, so a prime word like *departure* speeds responses to targets like *depart* when both prime and target words are presented in spoken form or written form, or when the prime is presented in one form and the target in another (Frost et al., 1997; Marslen-Wilson and Tyler, 1997; Marslen-Wilson et al., 1994).

*Masked priming* experiments also support a role for morphemes in lexical access. In masked priming studies, the prime word is presented visually followed by a pattern that covers the place where the prime was. Masking the prime stimulus prevents the visual system from taking up additional information about the prime word once the mask is displayed. Primes can be presented for very short amounts of time—as little as 43 ms, less than 1/20th of a second. When prime exposure duration is manipulated—some primes are shown for a very short time before being masked, some primes are shown for longer—different patterns of priming occur for semantic and morphological primes. At very short prime exposure durations, semantic priming (*doctor-nurse*) does not occur, but priming is very robust at longer prime exposure durations. The opposite pattern happens for morphological primes. At very short prime exposure durations, morphological priming (*apartment-apart*) is robust, but that priming effect disappears at longer prime exposure durations (Rastle et al., 2000). Results like these indicate that morphological priming effects, such as those that happen when *dishonest* is used to prime *honest*, do not reflect semantic overlap between the meaning of *dishonest* and the meaning of *honest*. Likewise, orthographic (letter) overlap does not account for priming in morphologically related pairs (if it did, *apartment* should prime *apart*, because all of the letters in *apart* are also in *apartment*). This suggests that morphological representations and processes play a unique role in lexical access, separate from semantic, phonological, and orthographic effects, as suggested by FOBS. Neuroimaging data also support a unique role for morphemes in lexical processing, because prime-target word pairs that share a root morpheme are associated with decreased neural activity in the left inferior frontal lobe, while other kinds of prime-target pairs are not (Bozic et al., 2007).

To summarize, the FOBS model proposes that word form representations are organized into bins. The set of bins is organized according to root frequency, and entries within each bin are organized according to surface frequency. This architecture explains why words with more frequent roots are processed faster than words with less frequent roots, and it can explain smaller effects of surface frequency. The model also explains why words that have pseudo-affixes are more difficult to process as a class than equally long and frequent words that have real affixes.

## Second-generation models

### TRACE

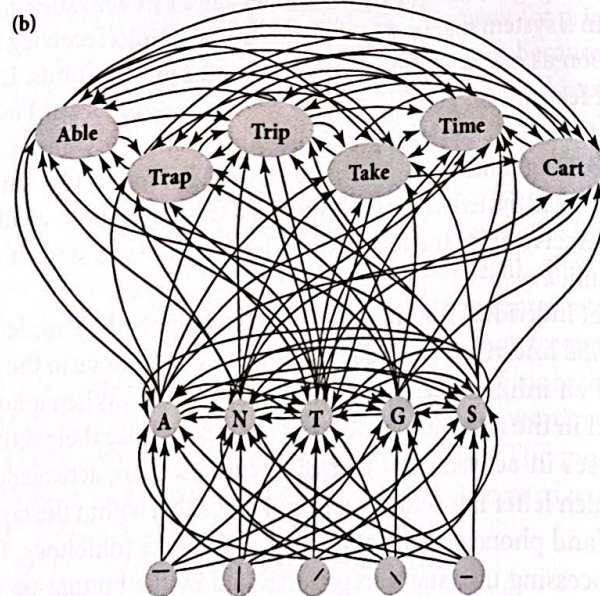
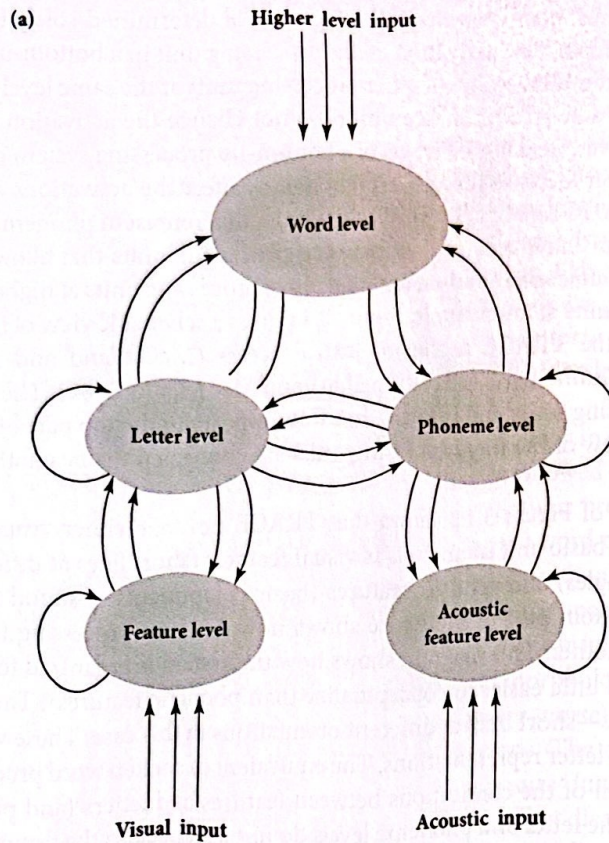
The TRACE model of lexical access differs from previous models in that, rather than having a serial, bottom-up, architecture, the model is highly interactive. In serial

bottom-up systems, activation of processing units is determined solely by stimulation provided by the input. The activation of one processing unit in a bottom-up system does not directly affect the activation of other processing units at the same level of the system. For example, an activated phoneme unit does not change the activation of other phoneme units. Activation at higher levels of a bottom-up processing system does *not* affect activation at lower levels of the system. Phonemes affect the activations of word units, but word units do *not* affect activation in the units that represent phonemes. Interactive processing systems have connections between processing units that allow units within the same level to affect one another, and that allow processing units at higher levels of the system to affect units at lower levels. Figure 3.11 gives a schematic view of the processing architecture in the TRACE model of lexical access (McClelland and Elman, 1986; McClelland and Rumelhart, 1981; Rumelhart and McClelland, 1982). The basic organization of processing units and information flow appears in the top part of the figure. A more detailed view of the way processing units are connected to one another appears at the bottom.


The top part of Figure 3.11 shows that TRACE can take either visual or auditory input. The most basic unit of analysis is visual features (short lines at different orientations, curves, angles) and acoustic features (basic components of sound in the speech stream). The bottom part of the figure shows how different processing units are connected to one another. This diagram shows how the system is organized for visual word processing (it's a little easier to conceptualize than phonetic features). The input to the system is *features*—short lines at different orientations in this case. These visual features are connected to letter representations. The equivalent in spoken word processing would be phonemes. All of the connections between features and letters (and phonemes) are excitatory, and the letter and phoneme levels do *not* feed back to the feature level.

The TRACE model assumes that activation is *cascaded*. Cascaded activation contrasts with *threshold activation*. In the logogen model, *threshold activation* means that a processing unit (e.g. a logogen) sits quietly until input causes its activation to exceed some threshold value. In a system that uses cascaded activation, units receiving input begin to send output as soon as any activation at all comes in from other units. Using cascaded activation, visual features in TRACE start to send activation forward as soon as they begin to be identified, so letter-level processing units start to become active soon after feature-level processing units start to become active. That means that letter representations start to become activated as soon as any visual feature has been identified, and you do not need to perceive all of the features of a letter before you start to activate letter-level processing units.

Notice also that individual features are connected to more than one letter-processing unit. The horizontal line visual feature has excitatory connections to the letters “A,” “T,” “G,” and “S,” and an inhibitory connection to the letter “N.” When a horizontal visual feature is detected in the input, all four of those letters increase their activation, and the letter “N” decreases in activation. If four different letters are activated, how does the system decide which letter it is actually seeing? Notice that within the layer of units representing letters (and phonemes), *all* of the connections are inhibitory. This means that when a letter-processing unit starts to get activated by the bottom-up input from the features, it will try to decrease the activation of the other letters that it is connected to. This pattern of connections leads to *lateral inhibition*—processing units within a layer of units in the network try to reduce or inhibit each other's levels of activation. This makes sense, because a feature can be only part of one letter. So, if the feature comes from the letter “t,” the representation of the letter “t” should try to inhibit other possibly competing candidates. After the bottom-up input has been received, inhibitory connections within a processing layer cause different letter representations to compete with one



**Figure 3.11** The TRACE model of lexical access. The top part shows the basic architecture. Connections with arrows at the end indicate excitatory influences; connections with round ends indicate inhibitory influences. *Source:* McClelland and Rumelhart (1981), American Psychological Association



**Figure 3.12** An example of degraded input that TRACE is good at processing

another, and the letter with the most support from the bottom-up features will eventually “win” the competition—its activation will increase and it will inhibit competing letter representations until eventually there is only one candidate left standing.

Letter representations have a slightly more complex relationship with word form representations. Letters have excitatory inputs to the words they are components of, and inhibitory connections to words that they are *not* components of. Activating the letter “A” will excite the words *able* and *trap*, and it will inhibit the word *time*. Letters have excitatory and inhibitory feedback connections from the word layer as well. This means that, as a word starts to become activated, it will feed excitation or inhibition back to the letter level. So, if the word *able* starts to get activation from the letter “A,” it will start to activate its other component letters, “B,” “L,” and “E,” via excitatory *top-down* feedback connections, possibly before those letter representations have been activated by bottom-up input. Simultaneously, activity at the word level for *able* will inhibit letter-level representations that are not present in *able*. This is one of the properties of TRACE that allow it to deal with degraded input, like that shown in Figure 3.12. A strictly bottom-up system would not be able to identify the right-most letter in Figure 3.12, because it could just as easily be an “R” as a “K,” and so a strictly bottom-up system might not be able to correctly identify the word as *work*. However, in the TRACE model, the intact letters “W,” “O,” and “R” would activate the word form representations “WORK,” “WORD,” and “WORM,” which would feed activation back to the letter level, and the combination of top-down and bottom-up activation from the remaining intact features would eventually cause activation of the “k” letter representation to exceed possible competitors.

TRACE also offers a good explanation of the *word superiority* effect. The word superiority effect refers to a class of behaviors indicating that we have an easier time recognizing and processing letters and phonemes when they appear in the context of a word than when they appear by themselves or in the context of a string of letters that does not make up a real word. The greater ease of processing letters and phonemes in the context of a word can be demonstrated in a number of different ways. In the 1800s, Erdmann and Dodge (Erdmann and Dodge, 1898, in Balota et al., 2006) showed that people could read a word containing up to 22 letters in the same amount of time it took them to identify 4 or 5 individual letters when those letters were not part of a real word. Other demonstrations of the word superiority effect come from phoneme and letter monitoring experiments (Foss and Swinney, 1973; Johnston and McClelland, 1973; Reicher, 1969; Savin and Bever, 1970; Wheeler, 1970; the phoneme restoration effects discussed in Chapter 2 also represent a form of word superiority).<sup>16</sup> In these experiments, subjects are given a target phoneme, like /s/, they listen to recorded speech, and they press a key as quickly as possible when they detect the presence of the target phoneme. Reaction times are faster when the target phonemes appear as part of a real word (and reaction times are also affected by how frequent the word-level representations are, suggesting that the word form representation is accessed *before* the phoneme is detected; Foss and Blank, 1980).

In other experiments, letters are presented either by themselves, as part of a non-pronounceable nonword (like *owrk*), or with the same letters rearranged to make up a real word (like *work*). The stimuli are flashed for a very brief amount of time, and then a test stimulus is presented consisting of two letters, *d* and *k*, for example. Subjects are asked to say which of the two letters appeared in the briefly presented stimulus. Notice that *d* and *k* can both be added to *wor* to make up a word, which eliminates guessing as a strategy for improving accuracy on the task when real words were presented. Despite this handicap, subjects were more accurate at identifying letters when the briefly presented stimulus was a word than in the other conditions. So, activating a word-level form representation helps people identify individual letters. This can be explained by the TRACE model by proposing that activation of word-level form representations strengthens the activation of letter-level representations via excitatory feedback as well as inhibition of possible competing letters that are not part of the activated word-level representation. McClelland and Rumelhart explain the word superiority effect on letter detection in this way (1981, p. 389): "the reason letters in words fare better than letters in nonwords is that they benefit from feedback that can drive them to higher activation levels."

To summarize the TRACE model: It is a highly interactive system. Bottom-up input, top-down feedback, and lateral inhibition combine to determine how much activation any given unit in the network enjoys. The TRACE model explains how and why we can deal with degraded input. The network computes the best fit to the degraded stimulus by simultaneously assessing multiple levels of representation, and a good fit at one level can compensate for a bad fit at another level. Finally, TRACE explains why letters in whole words are easier to perceive than individual letters by themselves. Feedback from the word layer boosts the activation of lower level letter representations.

### COHORT

The COHORT model is another prominent second-generation account of lexical access (Marslen-Wilson, 1987, 1989; Marslen-Wilson and Welsh, 1978). The COHORT model was developed specifically to explain lexical access for spoken words. It views the process of lexical access as involving three kinds of processes: *activation* (or *contact*), *selection*, and *integration*.

During the initial *activation* or *contact* phase of processing, multiple word form representations are activated in response to the auditory stimulus. COHORT views *contact* as being influenced only by bottom-up auditory information, and not by contextual information, and so activation in COHORT is referred to as an *autonomous* process—it is affected by auditory stimulation but not by other potentially relevant cognitive processes. As a result, stored representations of words that do not fit into the evolving context are activated anyway as long as they match the acoustic properties of the word stimulus.

*Selection* involves sorting through the activated word form representations to find the one that best matches the auditory stimulus. COHORT says that selection depends on the bottom-up stimulus, because bottom-up information activates word candidates, but it also depends on context. Words that fit better into the context have an advantage over words that do not fit, especially in cases where the bottom-up input is ambiguous between two or more stored word candidates.

*Integration* happens when the features of the selected word are incorporated into the evolving representation of the entire utterance. During integration, properties of the selected word—its grammatical class and meaning—are evaluated with respect to how well they fit with the preceding context.

Because COHORT deals with spoken input, it views lexical access and the activation of word form as resulting from a continuous evaluation of the similarity between the

auditory stimulus and stored word form representations. COHORT also views the process of lexical access as being radically incremental. Word representations are activated as soon as the initial sounds in the acoustic stimulus have been perceived, and it is possible for people to identify individual words in the speech stream *before* they hear the entire word.<sup>17</sup>

COHORT is called COHORT because the process of lexical access starts with a contact phase in which all words that match the perceived acoustic profile are activated. So within about 100–150 ms of the onset of a word, a whole group of matching candidate word forms become more available or accessible than usual. This group of activated word forms is called a *cohort*.<sup>18</sup> After this initial activation phase, the lexical access mechanism continues to check the list of activated candidates against further input from the speech stream, and it eliminates candidates that no longer match the input. Simultaneously, it checks the characteristics of each member of the activated cohort against requirements imposed by the context—the correct target word has to have the right syntactic and semantic properties to continue being activated.

COHORT makes very specific predictions about when, exactly, a word can be recognized and its meaning accessed. COHORT says that word recognition depends on reducing the set of activated words to the one that matches the acoustic input. The point where the COHORT is reduced to a sole survivor is called the *recognition point* (Marslen-Wilson, 1987). A word like *trespass* can be recognized well before the end of the word, because there are no other words besides *trespass* that are consistent with the onset *tresp* (Marslen-Wilson and Tyler, 1980).<sup>19</sup> COHORT allows for minor adjustments to the recognition point based on semantic or syntactic requirements imposed by context. Words that are highly predictable in context may be recognized a bit faster than less predictable words.

According to COHORT, word recognition is *contingent* on two factors: First, there has to be positive evidence for the presence of the word (e.g. the input *tres* provides clues that the word *trespass* is the matching word target). Second, the input has to rule out the presence of other words (e.g. the onset *tr* rules out the possibility that the matching word target is *tap*, *top*, *table*, or any other word that does not begin with *tr*). So, the word *trespass* can be recognized “early” (before the end of the word) only if lexical access includes a mechanism that knows when all other possibilities have been eliminated.

What evidence supports the prediction that multiple word candidates are activated early in lexical access, shortly after you hear a word’s onset? In a cross-modal priming experiment, participants listened to words like *captain* and *captive* that have the same onset (up to the *t* sound in the middle of the word). Visual probe words were presented either “early” in the word (at or before the *t* sound) or later in the word (during the final *-ain* or *-ive* syllable). When the probe word was presented early, two words were still compatible with the input so far (*capt* could continue and become either *captain* or *captive*). When the probe word was presented later, only one of the word candidates was still compatible with the auditory input. The target words were semantically related to one of the two meanings. The target word *ship* was presented, because it is related to *captain*. Alternatively, the target word *guard* was presented, because it is related to *captive*. The question was: Which target words would be primed by the auditory stimulus? If COHORT is correct, and word onsets activate all of the entries related to the auditory stimulus, then both *ship* and *guard* should be primed at the “early” probe point. At the “late” probe point, which comes after the recognition point, only one of the probe words should be primed. More specifically, priming should be observed only for the probe word whose meaning is associated with the surviving word candidate. This is exactly the pattern of priming that subjects experienced. Early in the word, both *ship* and *guard* were primed. Later on, only the matching meaning was primed (Marslen-Wilson and

Zwitserslood, 1989). So, if people heard *capti*, and the probe word appeared simultaneously with the *i* sound, only the target word *guard* was primed, and not *ship*. When these ambiguous spoken word onsets (*capt-*) were embedded in spoken sentences that made one meaning much more likely than the other, priming was still observed for both *ship* and *guard*, which suggests that context was not able to “turn off” or prevent access to the contextually inappropriate word (Zwitserslood, 1989). So, as Marslen-Wilson explains (1987, p. 89), “No contextual pre-selection is permitted, and context cannot prevent the accessing and activation of contextually inappropriate word candidates.”<sup>20</sup>

How does COHORT compare to TRACE? They differ with respect to how word form representations become activated. TRACE views word form activation as resulting from a process of competition and mutual inhibition. COHORT views word form activation as reflecting a massively parallel process without competition until the selection phase. The two accounts therefore make different predictions about what will happen as multiple word candidates become activated. According to TRACE, more activated word candidates are associated with less activation being gained by any one candidate, and greater competition between candidates. Because COHORT allows for unlimited parallel activation of word candidates, the number of activated candidates does not affect the speed with which the correct candidate is identified. To test this aspect of the models, Marslen-Wilson manipulated word onsets in a nonword detection task. Specifically, some of his stimuli were still consistent with many words at the point right before new auditory information rendered them nonwords. Other stimuli were consistent with very few words right before new auditory information rendered them nonwords. Presumably, to recognize that the stimulus is a nonword, people have to search through the set of activated candidate words to find a match. If words compete with one another, or if people search through the list in a serial fashion (as in FOBS), then the nonword judgments should take longer for bigger sets, and less time for smaller sets. However, reaction time data indicated that nonword judgments were made equally quickly, no matter how big the set of matching candidates was.

COHORT and TRACE also differ with respect to how similarity between the stimulus and stored word forms affects processing. TRACE relies on global similarity match to determine how active a stored word form becomes. So, it does not matter where a slight mismatch occurs in a word, at the beginning or the end. As long as the overall stimulus is close to the stored representation, the stored representation will become active.<sup>22</sup> In COHORT, word onsets are critical, because word onsets determine which representations will make it into the cohort, and which will be left out. As a result, mismatches at the beginnings of words should have greater effects than mismatches at the ends of words. According to TRACE, activation of word nodes will be a function of similarity (*bone* and *pone* will both lead to similar patterns of activation in the network). As a result, words that share offsets should prime each other’s meanings (because presenting *pone* activates the similar entry *bone*). The prediction, then, is that if you hear *pone*, you should respond faster to words associated with *bone*, like *arm*, *broken*, and *shin*. However, offset-matching primes are almost completely ineffective, suggesting that word onsets really do set the stage for lexical access (Marslen-Wilson and Zwitserslood, 1989; see also Allopenna et al., 1998; McQueen and Viebahn, 2007).

What evidence supports the psychological reality of recognition points? In one-syllable (*monosyllable*) words, the recognition point and the end of the word are one and the same, but for many multi-syllable (*polysyllable*) words, the recognition point comes well before the end of the word. The COHORT model says that words are recognized when the acoustic stimulus reaches the recognition point. So people can recognize words and access their meanings without having to wait until the very end of the word. Some experiments involve phoneme monitoring. Recall that phoneme monitoring speed is

affected by word frequency, suggesting that word identification precedes the monitoring response (Foss and Blank, 1980). It turns out that phoneme monitoring speed is strongly correlated with the recognition point. Words that have early recognition points lead to faster phoneme monitoring times than words that have later recognition points (Marslen-Wilson, 1984), whether the words are presented in isolation or as part of an extended utterance (Tyler and Wessels, 1983).

Nonword detection time also depends on when, exactly, the nonword stimulus diverges from real words that share the same onset. A nonword like *trenkitude* can be identified as a nonword faster than an equally long nonword like *cathedruke*, because *trenkitude* becomes a nonword sooner than *cathedruke*. The only English words with the same onset as *trenkitude* are *trench*, *trend*, and slight variations thereof (e.g. *trendy*), so *trenkitude* becomes a nonword at the *k*. *Cathedruke* has a potential word match, *cathedral*, up to the *dr*, and so becomes a nonword later. When people engage in nonword detection experiments, they respond faster to words like *trenkitude* than equally long nonwords like *cathedruke*, which provides further evidence for the special status afforded to recognition points. Both *cathedruke* and *trenkitude* provide bottom-up, positive evidence for real words, and so both should lead to roughly equivalent activation of word targets under both the logogen and FOBS models.<sup>21</sup> Plus, the FOBS and logogen processing systems might have to wait until the entire stimulus has been perceived before they can be certain that no logogen will fire or no bin will match the input. Thus, these models do not provide strong justification for why nonword detection times should differ between *cathedruke* and *trenkitude*.

The available evidence suggests that words are identified very quickly, and that the bottom-up information that is present at the point in time when words are identified often is not sufficient, by itself, to pick out one single word (Marslen-Wilson, 1987). One estimate is that 200 ms of spoken input is compatible, on average, with about 40 different words. This also is bad news for models, like FOBS, that say that word identification is based on an autonomous search process based on purely bottom-up information, because it appears that lexical access is well under way before a unique root morpheme can be identified. Another piece of bad news for FOBS is that lower frequency words affect recognition points as much as higher frequency words do. According to models like FOBS, higher frequency entries are searched before lower frequency entries are. So, if a word has a lower frequency competitor, this should not affect how quickly that word is accessed. However, COHORT makes a different prediction. According to COHORT, word candidates are activated based on their match to the initial sequence of sounds in the acoustic input, regardless of how frequent the word candidate is. So, if a higher frequency target word, like *rap*, has a lower frequency cohort member, like *rapture*, it will be recognized slower than an equivalent high-frequency word that does *not* have a lower frequency competitor. In fact, response times in a variety of tasks involving spoken words depend on the recognition point, independent of the frequency of the actual target word (Marslen-Wilson, 1987).

Activating multiple candidate words and continuously evaluating the goodness of fit between the acoustic input and the activated candidate set, as proposed by COHORT, confers a number of benefits to the listener. First, activating multiple candidates ensures that the correct word will be available for selection and further processing. Second, continuously evaluating the fit between the stimulus and the set of activated candidates ensures that the correct candidate will be selected as soon as the bottom-up input discriminates the correct target from similar alternatives. These characteristics therefore maximize the speed and accuracy of the lexical access mechanism.

The original version of the COHORT model did not have an explicit account of word frequency effects, but this shortcoming was addressed in follow-on versions of the

model. After it was established that frequency did affect word recognition times, independent of where the recognition point was, COHORT was modified to include different rise times in the level of activation for higher frequency and lower frequency word forms. This is essentially equivalent to the logogen model's move of lowering thresholds for higher frequency word forms. So the revised COHORT model, like the TRACE model, does not view word form activation as all-or-none. Instead, word forms can have no activation, a little activation, or lots of activation (consistent with behavioral evidence for early effects of frequency; Cleland et al., 2006; Dahan and Gaskell, 2007). For instance, if people are shown an array of objects, and they listen to a word that has an ambiguous onset, like *bell* (because the *be* part is consistent with a wide variety of words, *bell*, *bed*, *bet*, *bend*, etc.), they look more often and more quickly at a picture that goes with a high-frequency word than a picture that goes with a lower frequency word (Dahan, Magnuson, Tanenhaus et al., 2001). Responses to high-frequency visual word targets are also more facilitated by an auditory prime than low-frequency targets (e.g. the high frequency word *feel* is more primed when people hear the onset sounds *fee*; and the low-frequency word *robe* is less primed when people hear the onset sounds *roe*; Marslen-Wilson, 1990). To account for these kinds of effects, "Elements are not simply switched on or off as the sensory and contextual information accumulates, until a single candidate is left standing. Instead, the outcome and the timing of the recognition process will reflect the differential levels of activation of successful and unsuccessful candidates" (Marslen-Wilson, 1987, p. 93).

The revised COHORT model also alters its view of the input representation (Lahiri and Marslen-Wilson, 1991; Marslen-Wilson and Warren, 1994). The original versions of COHORT, logogen, FOBS, and TRACE all assume that a level of phonological processing units mediate between acoustic-phonetic features and word representations. That is, acoustic features activate phoneme nodes, and then phoneme nodes activate words. This information flow can produce catastrophic failure, however, if the bottom-up input is miscategorized, and the wrong phoneme is identified. To solve this problem, COHORT suggests that acoustic-phonetic features are directly connected to word-level representations (and that phoneme identification is a byproduct of activating word forms). That way, words that have similar acoustic-phonetic features will be activated. For example, *bat*, which has a voiced labial stop at its onset, would be partially activated when someone said *pat*, which has a de-voiced labial stop at its onset.

The direct mapping of phonetic features to word form representations also helps explain other *sublexical* (below the level of the word) effects on lexical access. Many English words contain *onset embedded* words. The word *lightning* starts with *light*. The word *hamster* starts with *ham*, which are words by themselves. However, it turns out that the string *ham* is pronounced slightly differently when it is a word all by itself compared to when it is just the first syllable of a bigger word. Specifically, *ham* has a longer duration, it sounds more like *haaaaam*, when it is spoken as an independent word (as in *This haaaaam tastes really good*) than when it is produced as part of a bigger word (as in *This hamster tastes really good*). These differences in pronunciation are detected by the auditory lexical access system fast enough to bias activation toward the matching word candidate. So the short word *ham* becomes more active when *ham* is pronounced with a longer duration (*haaaaam*); and the longer word *hamster* becomes more active when *ham* is pronounced with a shorter duration (Davis et al., 2002; Salverda et al., 2003, 2007; Shatzman and McQueen, 2006). Other sublexical properties, like where the stress occurs can also affect how rapidly individual word candidates become active (e.g. you pronounce *record* differently when it is a verb versus when it is a noun, but subtle differences in stress patterns also occur for stand-alone words like *ham* and the same segments that are embedded in larger words; McQueen et al., 1994).

DISTRIBUTED FEATURE MODELS

The parallel distributed processing enterprise continued to grow and develop with the invention of newer and more advanced mathematical models of lexical access. For example, Jeff Elman's *simple recurrent network* (SRN) model assumed that words were represented as a pattern of neural activity across a multi-layered network. As shown in Figure 3.13, the SRN model adapted TRACE's three-layered network and added to it a

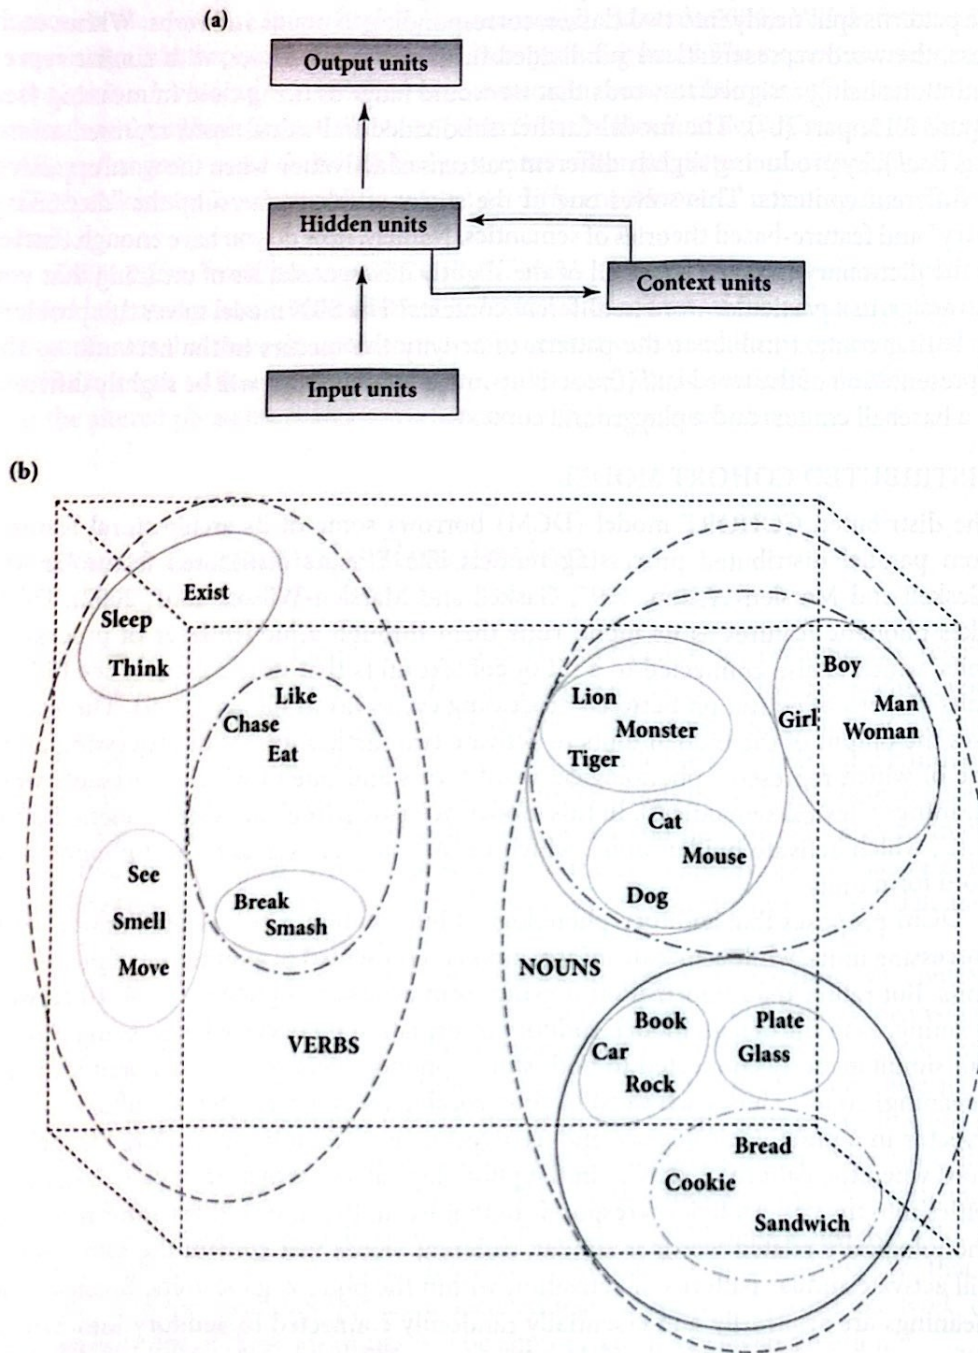


Figure 3.13 A schematic of Elman's simple recurrent network model of auditory word processing: (a) the architecture of the network; (b) the semantic space that emerged after the model was trained. Source: Elman, J. L. (2004), with permission of Elsevier

set of context units. The job of the context units was to store a copy of the activations in the hidden units between processing cycles. In this way, the network would respond not just to the current state of the input units, but also to recent events, as reflected in the activity of the context units. The explicit task that the network performed was to predict the upcoming word in an utterance. Before training, the network's connection weights were randomized, and then it processed a set of sentences one word at a time. As each word was encountered, the network tried to predict what the next word would be. When it made errors, the connection weights throughout the network were changed so that its output would more closely match the desired output the next time around. In this system, word identities can be represented as a pattern of activation among the hidden units. When Elman inspected these patterns after the network was trained, he found that the patterns split neatly into two classes, corresponding to nouns and verbs. Within each class, the word representations subdivided further into subclasses, with similar representations being assigned to words that we would judge as being close in meaning (see Figure 3.13, part (b)). The model further subdivided individual word representations (e.g. *book*), by producing slightly different patterns of activation when the word appeared in different contexts. This solves one of the sticky problems faced by the "dictionary entry" and feature-based theories of semantics. Namely, how do you have enough entries in the dictionary to take care of all of the slightly different shades of meaning that you can assign to a particular word in different contexts? The SRN model solves this problem by letting context influence the pattern of activity that occurs in the network, so the representation of the word *ball* (the activity in the hidden units) will be slightly different in a baseball context and a playground context.

#### DISTRIBUTED COHORT MODEL

The distributed COHORT model (DCM) borrows some of its architectural features from parallel distributed processing models like Elman's distributed feature model (Gaskell and Marslen-Wilson, 1997, Gaskell and Marslen-Wilson, 2001, 2002). DCM takes phonetic features as its input, runs them through a hidden layer of processing units, which is also connected to a set of context units that store a copy of the hidden units' pattern of activation between processing cycles (as in Elman's SRN). The system uses the output of the hidden units to activate two further groups of processing units, one of which represents phonological word forms and one of which represents word meanings ("lexical semantics"). In this model, acoustic stimuli activate phonetic feature units, which activate hidden units, which in turn activate semantic and phonological word form units.

DCM proposes that auditory/phonological information is represented in one set of processing units, while semantic information is represented in another set of processing units. But rather than information flowing from acoustic to phonetic to word form to meaning, as in the default model, auditory information is conceived of as being directly and simultaneously connected to both stored phonological codes and stored semantic (meaning) codes. Thus, each word in your vocabulary is represented simultaneously as a vector in a phonological space and as a vector in a semantic space. You recognize a word when the pattern of activity in the phonological and semantic units stabilizes and settles into the pattern that corresponds to that word. Because auditory information for phonologically related words is similar, different words that contain the same sounds will activate similar patterns of activation within the phonological units. Because word meanings are arbitrarily and essentially randomly connected to auditory information, different words with similar sounds will activate different, randomly assorted parts of the semantic space.

The simultaneous activation of phonological and semantic units has a number of consequences for word recognition and activation of meaning. In the initial moments of lexical access, when the onsets of words are heard, processing units in both the phonological and semantic spaces become activated. Activation in the phonological space will be coherent and mutually reinforcing because words with the same onset will share aspects of phonological representation. Activation in the semantic space will represent a blend of different semantic patterns. That is, the initial pattern of activation in the semantic network does not correspond to any of the stable states that represent individual word meanings. The pattern of activity in the semantic nodes therefore does not correspond to any familiar word meaning.

DCM is also good at explaining why words with multiple meanings are harder to process than words with multiple related senses (Rodd et al., 2004). Words with multiple meanings, like *bark* (as in *tree bark* and *Does your dog bark?* *The ancient mariner crossed the sea in a bark*) lead to less coherent activation in the semantic part of the network; and words with multiple senses, like *twist* (as in *Give the handle a twist*, *Can you do the twist?* *Oliver has gone round the twist*) lead to a more coherent pattern of activation. (Elman's SRN makes similar claims with respect to different flavors of meaning for words like *ball*).

DCM differs from the original COHORT model in that DCM places less emphasis on word beginnings as a critical element in lexical access. Part of the motivation for this is experiments that show that nonword primes can activate word form representations that differ in onset (*dob* and *tob* will both prime the word *bob*; Connine et al., 1993), as long as the altered phoneme shares some features with the original phoneme.

## Lexical Ambiguity Resolution

Many words have more than one meaning. The word *bank*, for example, can refer to a place where you keep your money or it can refer to a place next to a river where you go fishing. According to some estimates, over 40% of the words that you hear in English have more than one meaning (and this does not include the temporary ambiguities that happen when you hear words with onset-embedded words in them, like *ham* in *hamster*) (M. A. Gernsbacher, 1990). So what happens when you hear or read a word that has more than one meaning? Do you go straight to the contextually appropriate or correct meaning? Or do you have to sort through incorrect or contextually inappropriate meanings before you get to the correct one?

According to the *exclusive access* hypothesis, you can use cues from the context to immediately select the correct meaning of an ambiguous word like *bank*. When you hear or see the word *bank* you access only one meaning. If you are listening to a story about money, you access the financial institution meaning; and if the story is about fishing, you access the river-related meaning instead. But as we saw before, early events in word processing seem to involve activation of multiple candidates pretty much all the time. If visual and acoustic stimuli activate multiple word forms that they are associated with, maybe word forms simultaneously activate multiple meanings that they are associated with. This latter hypothesis is called the *exhaustive access* account. Exhaustive access says that you activate all of the meanings that are associated with an individual word like *bank*, even though only one of those meanings will be appropriate in any given situation.<sup>23</sup>

The exclusive and exhaustive access accounts were first tested in a series of priming experiments (Onifer and Swinney, 1981; Seidenberg et al., 1982; Swinney, 1979). In these

experiments, ambiguous words like *bug* were embedded in contexts that made one of their meanings more appropriate than the other. For example, the sentence might be, *The spy swept the room looking for concealed bugs*, in which case the “listening device” meaning of *bug* would be appropriate. In another case, the context sentence might be, *The cook picked up a bag of flour in the kitchen and saw the bugs*. In that case, the correct meaning would be the “insect” version. To assess which meanings subjects accessed, their responses to words associated with the different meanings were measured. If people access the “listening device” meaning of *bug*, then they should respond faster to the target word *listen* than to an unrelated control word. If people access the “insect” meaning of *bug*, then they should respond faster to the target word *insect* than to an unrelated control word. So after people heard the word *bugs*, they responded as quickly as possible to a test word flashed up on a computer screen. The test word could be related to one or the other meaning of the word *bug*, or it could be unrelated. The difference in response time between related and control words provides an index of how activated the related word meanings were. Hearing the word *bugs* made people respond faster to target words related to either of its meanings, no matter which meaning was appropriate in context. These results are more compatible with the exhaustive access hypothesis, and they are incompatible with the exclusive access hypothesis. People do not appear to select only the right meaning. Both contextually appropriate meanings and inappropriate meanings are activated when people hear an ambiguous word like *bugs*.

If appropriate and inappropriate meanings are both activated when we hear an ambiguous word, how do we ever figure out the correct meaning of an utterance? If the “insect” meaning gets activated in a “listening device” context, why don’t we interpret the utterance as referring to insects rather than listening devices? The answer is that context does affect meaning selection eventually, even though it does not appear to prevent incorrect meanings from being activated in the first place. In follow-on experiments investigating meaning selection for ambiguous words, experimenters manipulated the amount of time that elapsed between the ambiguous word and presentation of the target word. The amount of time that passes between presentation of the ambiguous word and presentation of the target is called *stimulus offset asynchrony* (or SOA). In some studies, target words are presented immediately after the ambiguous word in some conditions, and they are presented at longer SOAs in other conditions. Different patterns of results are observed at different SOAs in experiments looking at ambiguous word processing. If target words are presented immediately after the ambiguous word, all of a word’s associated meanings are primed. But if you wait until 250–500 ms after the ambiguous word to present the target word, you get a different pattern of results. At longer SOAs, only meanings that are appropriate in context are primed. This means that, although all of the meanings of *bugs* are activated when you hear the word, context causes you to deactivate or suppress the inappropriate meaning after a short period of time. Thus, your long-term representation for the utterance will contain only the appropriate meanings, and your interpretation will not be cluttered with inappropriate meanings.

### *Does context influence meaning selection for ambiguous words?*

To explain how context influences meaning selection in ambiguous word processing, we need to introduce the idea of *meaning dominance*. Many words have multiple meanings but those meanings are not all created equal. Some meanings occur more frequently than others. For example, the “metal ore” meaning of *tin* (as in *This can is made out of tin*)

*tin*) is far more frequent in American English than the “container” meaning of *tin* (as in *I bought a tin of beans*). So, *tin* is like *bugs*, in that they both have more than one meaning. But *tin* is unlike *bugs* in that one of its meanings occurs more often than the other. This property of ambiguous words is referred to as *meaning dominance*. Some ambiguous words have one frequent (*dominant*) meaning, and other less frequent (*subordinate*) meanings. Let’s call this kind of word a *biased* ambiguous word. Other ambiguous words have two roughly equally frequent meanings. Let’s call this kind of word a *balanced* ambiguous word. It turns out that different kinds of ambiguous words, biased and balanced, have different effects on people’s behavior, and these differences reflect different underlying meaning-access processes.

When balanced ambiguous words are presented in a neutral context, they behave like *bugs* did—both meanings are activated simultaneously to roughly the same extent. That can be demonstrated in experiments involving eye tracking. In an eye-tracking experiment, people read texts—sentences, in this case—and their eye movements are recorded. Because eye movements are linked to the mental processes involved in interpreting the text (Rayner, 1998; Rayner and Pollatsek, 1989, 2006), we can estimate how much difficulty people have interpreting a given piece of text by measuring how long they look at that piece of text. When people read balanced ambiguous words in a sentence, *The woman saw the bugs ...* they fixate those words longer than matched control words that have only one meaning. When people read biased words, like *tin*, a different pattern emerges. People read biased words just as quickly as matched unambiguous control words, suggesting that they are only activating one meaning. Note that in these cases, the context that comes before the balanced ambiguous word does not indicate which meaning is appropriate. These are called *neutral* contexts. But we can change the context so that it favors one or the other of the word’s meanings. For example, we could change the context so that it favors the “insect” meaning of *bugs*, as in *What crawled out from under the sink was a bunch of bugs*.

This kind of biasing context has different effects depending on whether the critical word is balanced (*bugs*) or biased (*tin*). Biasing contexts cause balanced ambiguous words to be processed as quickly as matched unambiguous words. So you read *bugs* just as quickly as a word with only one meaning when the preceding context points you toward one of its meanings (Duffy et al., 1989; Rayner and Duffy, 1987; see also Tabossi et al., 1987; Tabossi and Zardon, 1993, who found a similar pattern of results in a set of priming experiments).

Biasing contexts have different effects on biased ambiguous words, depending on whether the dominant or subordinate meaning is appropriate. If the context makes the dominant meaning appropriate (e.g. *The miners dug into the mountain looking for tin*), the ambiguous word is processed just as quickly as a matched control word that has only one meaning (e.g. *The miners dug into the mountain looking for gold*). But if the biasing context points toward the infrequent meaning of a biased ambiguous word (*The miners went to the store and saw that they had beans in a tin*), it takes people a long time to read the word *tin*, suggesting that they are having a hard time accessing its subordinate (less frequent) meaning. Neuroimaging also shows that the same factors of meaning dominance (balanced versus biased) and context (supporting the dominant or subordinate meaning of a biased ambiguous word) affect the neural response to sentences containing ambiguous words (Mason and Just, 2007).

Balanced ambiguous words are read slowly in neutral contexts (suggesting exhaustive access to all of their meanings) and quickly in biasing contexts (suggesting that biasing context helps people select an appropriate meaning of a balanced ambiguous word). Biased ambiguous words are read quickly in neutral contexts (suggesting rapid access to one meaning), quickly when biasing context points toward the dominant meaning, and

slowly when biasing context points toward the subordinate meaning. This pattern of response inspired the *reordered access theory* (Duffy et al., 1989; Rayner and Duffy, 1987). According to *reordered access*, access to word meanings is influenced by two interacting factors. The first factor is meaning dominance—more frequent meanings will be easier to access than infrequent meanings. When you encounter a word, the bottom-up input activates all of the semantic representations associated with the word. Word representations are organized as in the TRACE model, so that when more than one representation is activated by a word, the activated representations compete with one another. Biased ambiguous words are easy to process because the dominant meaning wins the competition quickly. Balanced ambiguous words are more difficult to process because the two competing representations are more evenly matched, and it takes longer for competition to select a winner. The second factor that influences meaning selection is the context that a word appears in. When context and meaning dominance both favor the frequent meaning of an ambiguous word, competition between multiple activated word meanings is short-lived—the dominant meaning wins the competition very quickly. When context favors the less frequent meaning, its activation is raised to the point where it becomes an effective competitor with the more dominant meaning. As a result, the subordinate meaning can be selected when context favors it, but it takes more time for the subordinate meaning to beat down the more frequent dominant meaning.

## The Neural Basis of Lexical Representation and Lexical Access

Investigating what happens when people experience brain damage (*neuropsychological* approaches) and measuring activity in the intact brain (*neurophysiological* and *neuroimaging* approaches) are great ways to study how word meanings are organized in the brain and how the brain performs the processes required for lexical access. Neuropsychological approaches have demonstrated that knowledge of concepts and knowledge about word forms are handled by quasi-independent systems in the brain. That is, people can have intact knowledge of concepts, without being able to recover information about the word forms that refer to those concepts and vice versa (e.g. Damasio et al., 1996; Tranel et al., 1997). Neuroimaging experiments support a shared semantic system for words and pictures, but some brain areas respond more to words than pictures, and vice versa (Vandenberghe et al., 1996; Wagner et al., 1997). When subjects judged the similarity between word meanings or pictures, both kinds of stimuli activated a network of left hemisphere brain areas including the superior occipital cortex, the inferior (bottom) temporal lobes, and the inferior frontal lobes (see Plate 3, top). Word-specific activity was observed in a region of the left hemisphere in the superior temporal and middle (toward the center of the brain) anterior (front) temporal lobes (Plate 3, middle), as well as in the frontal cortex. Pictures selectively activated a region near the left superior temporal sulcus (Plate 3, bottom).


In the normally functioning brain, nonlinguistic conceptual knowledge and linguistic knowledge about words somehow combine to produce meaning when words are heard or read. How is this done? And where is it done? Answering where word meanings are stored in the brain and how the brain activates those meanings in response to auditory and visual stimulation runs into immediate complications when you consider that there are many different ways to classify words at many different levels of abstraction (open class vs. closed class, noun vs. verb, animate vs. inanimate, regular vs. exception, his-

frequency vs. low frequency, animal vs. vegetable vs. mineral, and on and on and on). So you should not be surprised to learn that, although widespread left-lateralized brain activity occurs when people listen to or read words, the specific pattern of brain activity reflects an interaction of word and task properties (Booth et al., 2003; Friederici et al., 2000; Posner and Raichle, 1994). Brain responses that depend on word properties can be observed in both neuropsychological and neuroimaging studies. Some aphasic patients appear to have greater difficulty retrieving information about verbs than about nouns, and others have the opposite problem (Caramazza and Hillis, 1991; Damasio and Tranel, 1993). These differences in the ability to retrieve words happen even when the two versions are nearly identical, as in the noun-verb pair *a crack* and *to crack*, and may be more severe for words that are less semantically complex, despite their greater frequency in the language (Breedin et al., 1998).

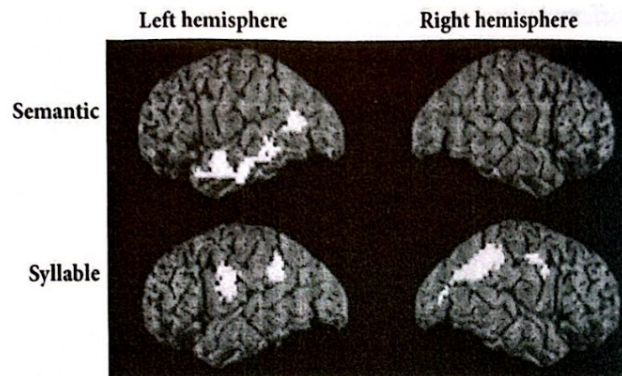
Different kinds of words appear to activate different brain areas, potentially reflecting differences in the way the brain represents the concepts to which the words refer. In a landmark PET study, Alex Martin and colleagues showed pictures of animals and tools and had participants say the names of the pictured object silently to themselves (Martin et al., 1996).<sup>24</sup> Different patterns of activation in the brain were observed for animals and tools. Greater activity was observed in occipital regions when naming animals, and greater activity was observed in inferior frontal regions when naming tools (see Plate 4).<sup>25</sup> Other ERP and imaging studies have shown that words referring to *concrete* entities (like *cat*, *dog*, and *table*) produce different patterns of neural response than function words (like *between*, *because*, and *where*) (Nobre and McCarthy, 1994; Nobre et al., 1997).

Neuroimaging and neuropsychological studies also show that the brain areas involved in processing a word differ depending on what kind of task people are doing when they encounter the word. When people are asked to generate the action that goes with a noun like *hammer*, activity is focused in the anterior cingulate gyrus, the left inferior frontal lobes, and the right cerebellum (Petersen et al., 1989; Posner et al., 1988; Posner and Raichle, 1994). In seminal PET imaging studies, Mike Posner and colleagues measured the brain's response to sets of nouns under different task conditions that he hoped would engage different brain regions. In one condition, brain activity during passive perception of words was compared to a fixation-cross baseline (that is, subjects just looked at a "+" on the screen during the baseline task). In the *dangerous animals* condition, participants would view a list of nouns and decide whether each one represented a dangerous animal or not (this is a type of *semantic categorization* task). In the *action generation* task, participants viewed each noun (e.g. *hammer*) and said an action that a person would undertake with that object (e.g. *pound*). Passively viewing words and not doing anything with them led to greater activity in the occipital lobes in both cerebral hemispheres. Tasks that tapped semantic features (dangerousness) or associations (between nouns and actions) produced left-lateralized activation in the frontal lobes (see Figure 3.14).

Different patterns of brain activity are also observed between tasks that focus on the semantic properties of words in contrast to their phonological properties. PET data showed significant neural activity throughout substantial parts of the left temporal lobe in response to semantic judgments, and bilateral (both sides of the brain) activation in more dorsal (toward the top) areas in response to judgments about how words sound (Price et al., 1997; see Figure 3.15). Other neuroimaging experiments also indicate that left prefrontal involvement in semantic processing differs across tasks involving the same words (Goldberg et al., 2007). Questions that tapped abstract, verbally acquired knowledge about animals led to stronger activation of left frontal regions than questions that tapped more concrete, directly observable properties of animals, even though the different kinds of questions were equally difficult to answer (see Plate 5; see also



**Figure 3.14** Results from PET neuroimaging experiments. Triangles indicate greater neural activity when participants passively looked at words, compared to a fixation cross baseline condition (solid black shapes indicate left-hemisphere activity, open shapes indicate right-hemisphere activity). Squares indicate areas with greater activity in the action-generation task versus repeating nouns out loud. Circles indicate areas with greater activity during the dangerous animals task than passive viewing of nouns. *Source:* Posner et al. (1988), American Association for the Advancement of Science



**Figure 3.15** PET data showing the neural response to a semantic judgment task (top) and a phonological judgment task (bottom). *Source:* Price et al. (1997), The MIT Press

Bright et al., 2006; Demb et al., 1995).<sup>26</sup> Different regions of the left prefrontal cortex also appear to be involved in tasks that tap semantic versus phonological knowledge associated with individual words. Some subregions are more activated when participants judge whether two words have similar meanings, while other subregions are more activated when participants judge whether two words rhyme (Heim et al., 2005; Roskies et al., 2001; see also Mainy et al., 2008). TMS results also suggest an anterior-semantic, posterior-phonological organization of the left inferior frontal region (Gough et al., 2005). Other studies have shown differences in brain activity in the left inferior frontal lobe occurring in response to lexical decision and verb-generation tasks (Frith et al., 1991).


Keeping in mind that the precise pattern of brain activity that is associated with a word depends on characteristics of the individual word and the task that the person is engaged in, word processing tasks generally activate a network of left-hemisphere regions. Right-hemisphere activity is also seen in some circumstances, especially for processing of words referring to abstract concepts (Kiehl et al., 1999), but greater neural activity in word processing tasks normally occurs in the left hemisphere.

Processing models assume a separate set of input representations for auditory and visual word processing (e.g. Coltheart et al., 2001; McClelland and Elman, 1986; McClelland and Rumelhart, 1981), and this division is reflected in different patterns of

activity in spoken and visual word processing. Auditory input more strongly activates Wernicke's area (near the junction of the occipital, temporal, and parietal lobes); and visual input may not activate this area at all (Howard et al., 1992; Petersen et al., 1988). Brain regions involved in auditory word processing include the superior temporal lobes bilaterally. These regions are involved in analyzing the acoustic and phonetic properties of the input (Kluender and Kiefte, 2006; Scott et al., 2000). Some theorists suggest that a portion of the superior (top) posterior (toward the back) temporal lobe in the left hemisphere contains a *phonological word form area* that is responsible for mapping acoustic information onto stored representations of individual words (e.g. Friederici, 2002).

Basic visual processing of written words is conducted by portions of the *striate* ("stripey") and *extrastriate visual cortex* in the occipital lobes in both hemispheres (these areas also respond to other complex visual stimuli). Further input processing of written words is associated with activity in the *visual word form area*, an area in the left hemisphere anterior (toward the front) to basic visual processing areas that are near other *perisylvian* cortical regions that are thought to be involved in phonological and semantic processes (Cohen et al., 2002; Dehaene et al., 2002; McCandliss et al., 2003; Nobre et al., 1994). This area responds to pronounceable letter strings, but not to spoken words or word-like stimuli; and it does not respond to complex visual stimuli other than words. Figure 3.16 displays the location of the visual word form area.<sup>27</sup>

When a stimulus has activated auditory or visual input codes, additional neural activity will reflect access to aspects of the words' semantic and syntactic properties, and the integration of these features into a representation of the ongoing discourse. At some point, this activity will not depend on whether the word was heard or read, and so it will reflect modality-independent information associated with the word in question. Such *post-lexical* processes are associated with widespread left-lateralized neural activity spread across regions in the anterior occipital cortex forward through inferior parietal lobes, the medial and inferior temporal lobes, the temporal poles, and the inferior frontal lobes (Friederici et al., 2000; Howard et al., 1992). However, only a small



**Figure 3.16** The visual word form area. The left hemisphere appears on the right side of the figure. *Source:* Cohen et al. (2002), Oxford University Press

portion of this activity appears to be task independent. Activity in the inferior (bottom) temporal lobe and frontal lobes appears to increase when tasks focus on semantic (as opposed to syntactic or visual features) of words. Activity in the superior temporal lobes appears to be more related to phonological analysis, which has led some theorists to propose a neural organization scheme in which dorsal (toward the top) brain areas are involved in phonological and motor analysis of speech input, while middle areas are involved in syntactic-relational information, and ventral (toward the bottom) areas are involved in retrieval of semantic information (Shalom and Poeppel, 2008; Thompson-Schill et al., 1999).

Similarly, the brain may be organized along the posterior-anterior dimension, with more posterior regions involved in retrieval of more basic features and more anterior areas involved in processing complex combinations of features and other relational information (Noppeney et al., 2007; Randall et al., 2004). This approach is supported by neuropsychological studies showing dissociations between verbs and nouns. Problems dealing with verbs are more associated with frontal lobe damage, while problems dealing with nouns referring to concrete entities are associated more with temporal lobe damage (Caramazza and Hillis, 1991). An alternative hypothesis proposes that the posterior-anterior organization may reflect a functional-perceptual distinction, as stimuli relating to tools (defined by their functions) appear to activate posterior brain regions more strongly than stimuli relating to animals (which are distinguished more by what they look like than what they're good for; Tranel et al., 2005), whether those stimuli are conveyed as words, pictures, or as sounds associated with objects (like *moo* for cow, or a snipping sound for scissors).

Hanna Damasio and her colleagues tested over 100 brain-damaged patients and correlated their performance on tasks involving the names of tools, animals, and people (Damasio et al., 1996). By mapping the locations of brain lesions and comparing lesion location with performance for different kinds of words, Damasio's group found that brain damage in posterior areas of the left temporal lobe correlated with deficits on tools, damage to adjacent more anterior regions was correlated with deficits on animals, and damage to the temporal pole was correlated with deficits on people (see Plate 6). These data could be interpreted as showing that different concepts are represented by different underlying neural systems. But critically, the vast majority of Damasio's patients could define concepts that they could not name. So, a patient might respond to a picture of a skunk by saying, "Oh, that animal makes a terrible smell if you get too close to it; it is black and white, and gets squashed on the road by cars sometimes" (Damasio et al., 1996, p. 499). So the patient's problem is not that they lack knowledge about the concept; rather there is something that prevents them from coming up with the name even though they can access aspects of the concept's meaning. As a result, Damasio and colleagues suggest that the temporal regions affected by their patients' lesions are responsible for *intermediary processes* that provide the links between distributed conceptual knowledge and phonological word form knowledge that is supported by language areas in the superior temporal lobe and the temporal-parietal-occipital junction. That is, the different regions of the temporal lobe are not storing localized conceptual representations. Instead:

*when the concept of a given tool is evoked (based on the activation of several regions which support pertinent conceptual knowledge) ... an intermediary region becomes active and promotes (in the appropriate sensorimotor structures) the explicit representation of phonemic knowledge pertaining to the word form which denotes the given tool. When a concept from another category is evoked, that of a particular person for example, a different intermediary region is engaged. (Damasio et al., 1996, pp. 503-504)*

## How are word meanings represented in the brain?

One of the enduring controversies in language science relates to how word meanings are represented in the brain, and a good way to get into this debate is to look at the phenomenon of *category-specific* semantic deficits. Category-specific semantic deficits happen when an individual has difficulty understanding the meanings of some types of words but not others. In particular, there seems to be a distinction between the processing of words that refer to natural kinds (animals, plants, and foods) and artificial or man-made objects (tools, buildings, and objects). The existence of category-specific deficits has been used to argue for localized semantic representations or separate semantic systems for living and nonliving things (e.g. Caramazza and Hillis, 1991; Pinker, 1994). According to the *localizationist* theories, semantic memory has been divided into separate categories by natural selection, because those categories represent biologically important domains. That conceptual division is reflected in a physical division of different kinds of concepts in different physical locations in the brain. So, if a lesion strikes the area that is responsible for representing conceptual knowledge of tools, an individual with that kind of damage will not be able to comprehend or produce words relating to those lost concepts. Other concepts may be completely spared, however, because they are physically instantiated in an undamaged region of the brain.

The localizationist approach contrasts with the distributed representation approach. According to the distributed representation approach, concepts are represented as coordinated patterns of activity across a wide variety of brain regions. In this kind of account, word representations can be thought of as a kind of *Hebbian cell assembly* (e.g. Pulvermüller, 1999). Hebb was a theorist who was active in memory research at the dawn of the cognitive revolution. He argued that concepts (and other kinds of long-term memories) consisted of linked groups of neurons. Groups of neurons are tied together with excitatory connections so that any time one of the members of the group becomes active, all of the other members of the group also become active. In this way, a simple retrieval cue could activate a rich and complex array of knowledge. You can think of a word as a retrieval cue that activates a subassembly representing the word's form, and the concepts and associations that become activated when you hear the word reflect the other components of a Hebbian cell assembly. How do Hebbian cell assemblies form in the brain? According to Pulvermüller (1999), such assemblies form when different groups of neurons are active at the same time. For word learning, this happens when neurons that respond to the sound of a word fire at the same time as other neurons that are responsible for representing perceptual (visual, tactile, auditory, etc.) and functional (what do you do with the object?) properties of the object. Once these associations are formed, you can access the sound when the perceptual and functional properties are activated (by direct experience or recollection); and the sound will similarly activate perceptual and functional representations associated with the name (as in the embodied semantics account). The fundamental claim is that word representations reflect neurally distributed groups of neurons that fire together when one subcomponent of the cell assembly becomes activated.

The idea that different kinds of words are represented in different parts of the brain has been investigated by looking at how word knowledge breaks down following brain damage and by neurophysiological and neuroimaging studies of normally functioning individuals. At first blush, the existence of category-specific deficits in semantic knowledge and word processing would seem to favor localization over distributed representations. Localized representations offer a quick and efficient explanation for why one category would go away but others would not, and one of the hallmarks of distributed

systems is graceful and gradual reduction in function following damage. However, a detailed look at the available evidence and a fresh look at the organization of semantic memory provides a major boost to the distributed representations position.

First, consider that the loss of knowledge of living things is more common than loss of knowledge of artificial kinds. Better preservation of knowledge about artificial kinds than natural kinds can be demonstrated in *confrontation naming* (patients try to say the word that goes with a pictured animal or object), *category fluency* (patients try to give as many examples as possible of a given category, like *plant*, or *animal*), and *definition tasks* (patients try to give a definition for a word). The box below provides some example definition task responses from patients with category-specific deficits for living things. These patients typically have damage to inferior and anterior portions of the temporal lobes (Saffran and Schwartz, 1994). Deficits for tools are associated with damage to posterior portions of the temporal lobes, and the portion of the parietal lobe near the occipital-temporal junction (Damasio et al., 1996). Worse performance for nonliving than for living categories is generally observed only in patients with the most severe semantic deficits (Moss et al., 1998). The localizationist/separate systems position explains why knowledge of living and nonliving things can differ, but it does not explain why deficits occur for living things more often than nonliving things.

Second, consider that the degradation of semantic knowledge is not all-or-nothing. Some information about the impaired category is preserved, and patients do better on some tasks than others, depending on how much detailed knowledge is required to do the task, independent of whether the task taps into knowledge of living or nonliving concepts. Bright and his colleagues (Bright et al., 2006) used a technique similar to Bates and colleagues' voxel-based lesion symptom mapping technique (VLSM) to investigate the relationship between conceptual knowledge, word processing, and the brain. In this study, patients with brain damage in different parts of the brain performed tasks that involved different types of words—natural kinds (like *cat*, *horse*) and artifacts (like *hammer*, *automobile*). The researchers measured the neural response to different kinds of objects and different kinds of information-processing tasks using fMRI. As in VLSM,

### DEFINITIONS OF LIVING AND NONLIVING THINGS PROVIDED BY PATIENTS WITH A CATEGORY-SPECIFIC DEFICIT FOR LIVING THINGS

*Patient RC*

*Bee*—"Bees are animals. And I've forgotten what they look like. But they're two-eyed, similar to humans. Two eyes of a see-through. Or a hearing, of—two ears. Of a mouth—of an eating, drinking."

*Bike*—"Bikes are two-wheeled, some are four-wheeled—of a learning of, of a learning for children ... or a two in the centr-ish, in the two on either side on the back, of a balance, of a

get a go, of a p ... of a, I don't know what it's called ... of a pedalling, of a pedalling and a steering and a four-wheeled as a start of a learn."

(Moss et al., 1998, p. 304)

*Patient JBR*

*Snail*—"An insect animal."

*Briefcase*—"Small case used by students to carry papers."

(Warrington and Shallice, 1984; in Saffran and Schwartz, 1994, pp. 513–514)



the researchers measured where the peak response in the brain occurred, they assessed how well patients did on the different kinds of objects and tasks, and they correlated the neural response with accuracy on the different tasks. Patients with greater signal intensity in the anterior (front) part of the temporal lobes did better on tasks involving natural kinds, when those tasks called for judgments about the fine details that you would need to know to discriminate between different concepts (e.g. *Do cats have whiskers? Do dogs bark?*). Questions that tapped shared features (*Do cats have legs? Do dogs have fur?*) were *not* associated with greater signal intensity in the anterior temporal lobes. If knowledge of natural kinds was supported in general by a neural network located in the anterior temporal lobes, both kinds of questions should have led to similar signal intensities in that brain region. Thus, these results are more compatible with a distributed account of semantic knowledge, with increasingly complex features and combinations of features supported by more anterior regions, but without a dissociation between living and non-living categories in terms of where in the brain associated information is stored.

Thomas Grabowski and colleagues' (2001) PET neuroimaging study involving famous landmarks and people also creates problems for the localizationist account of semantic representation. According to the localizationist account, concepts from different categories (e.g. animals, tools) are represented in different regions of the brain and accessed by different neural systems. The perceptual–functional approach argues instead that left hemisphere semantic processing regions are organized along the posterior–anterior axis such that functional and more general features are represented more posteriorly, and more complex combinations of features are represented more anteriorly. According to the localizationist account, pictures of landmarks (e.g. The Washington Monument, The Hubert H. Humphrey Metrodome, Carhenge) and people should activate different brain regions than pictures of famous people. According to the perceptual–functional approach, discriminating between landmarks and people both involve assessing fine-grained details, and so discriminating landmarks and people should both depend on more anterior regions, such as the temporal pole. In Grabowski and colleagues' PET study, unique landmarks and famous people both activated the left temporal pole and no differences in neural activity in any brain region were found between the famous landmarks and famous people conditions. These data are straightforwardly compatible with the perceptual–functional approach, but they pose problems for accounts that propose separate localized representations for living and nonliving categories. Other neuroimaging studies have also shown that the same brain regions become activated by words in different conceptual categories (Chao et al., 2002).

Category-specific deficits can be explained in a localizationist framework by proposing that certain concepts are represented at particular places in the brain, and the semantic system is organized so that similar concepts are represented in nearby locations in the brain. So a lesion that wipes out the “cat” concept is likely to wipe out similar concepts as well, but may spare semantically dissimilar concepts. The *correlated features approach* makes different representational assumptions, and offers a different way to explain category-specific deficits. According to the correlated features approach, semantic/conceptual knowledge is represented in distributed neural networks. Because semantic representations are distributed, you can't point to a place in the brain and say, “That is where the concept ‘cat’ is stored.”

The assumption of distributed knowledge has two major consequences. First, when you hear the word *cat* or think about cats, a wide variety of brain regions become activated, each of which may be responding to different aspects of the meaning of *cat*. This approach is similar to Pulvermüller's (1999) cell assemblies approach and other distributed representation and processing theories. Second, the same large, distributed network of brain regions is responsible for all of our semantic/conceptual knowledge. So,

knowledge about cats and other natural kinds is stored and activated by the same distributed system that is responsible for our knowledge about tools and other non-natural kinds. But if knowledge about animals and tools is spread all over the brain, and if knowledge about cats is handled by the same system that handles knowledge about hammers, how can we have a problem with just animals or with just tools?

The answer may be in the structure of the concepts themselves (Moss et al., 1998). Concepts consist of different kinds of features. Some features are *correlated* and some are *distinctive*. Correlated features are shared by many individual examples within a category. Distinctive features are those properties that make the difference between being one thing and being another. Living things have properties that tend to be highly correlated, and differences between different kinds of living things depend on minor differences in very specific (distinctive) attributes. As a result, if you know one thing about an animal, a lot of other properties are highly likely. If you know that something has eyes, it's almost certain that it has a mouth, a nose, lungs, four limbs, ears, and so forth. To tell the difference between different animals, you need detailed knowledge about subsets of properties. Does it have stripes and whiskers? Then it might be a tiger. Does it have stripes and a mane? Then it might be a zebra. By contrast, nonliving things are more likely to have un-correlated properties—knowing one thing about a nonliving thing does not make prediction of its other properties very easy. Nonliving things are also more likely to have multiple distinguishing features than living things. If you know an object has a handle, that does not allow you to predict whether it will have a bowl at the end of the handle, a flat head, or a point. But if you know that the rest of the object has a sharp edge, you are very likely to be dealing with a knife and not a hammer. Patients with category-specific deficits have more trouble with properties that discriminate between concepts that have many correlated features; and they have little trouble dealing with *common* features that occur across many different examples within a category. If someone has a specific deficit about animals, and you ask them about the properties that animals have in common (eyes, ears, legs, etc.), their performance is normal or near normal. If instead, you ask about *distinctive features*, they have a big problem. Patients like RC (reported in Moss et al., 1998; see also Bunn et al., 1998) could provide numerous shared features of animals, but not for artifacts. For artifacts, RC was able to provide distinctive features but not shared properties.

Neuroimaging data from patients with a category-specific deficit reinforce the idea that these deficits result from a general inability to deal with distinctive features generally, rather than a particular kind of concept (e.g. living vs. nonliving). In Peter Bright and colleagues' (2006; see also Devlin et al., 2002) study, fMRI was used to image brain activity in patients with category-specific deficits. While they were being scanned, they answered questions about pictures of living and nonliving objects. The nonliving objects included vehicles, which are an interesting case, because they have many correlated features (e.g. engine, steering wheel, tires, seats, and so forth), and the features that distinguish them tend to be highly idiosyncratic (all sedans look alike to someone who drives a truck and vice versa). So in terms of feature structure, vehicles are a lot like animals. In the fMRI experiment, some of the questions asked about shared properties (*Does it have tires? Does it have eyes?*) and some asked about distinctive properties (*Does it have claws? Does it have a peace-symbol on the hood?*). The idea was to find out if the patients had trouble with living versus nonliving things and whether brain activity and question responses differed between animals and vehicles. If instead the feature structure drives subjects' performance, then the patients should do worse on distinctive feature questions than on shared feature questions, whether the targets were living or nonliving. The fMRI results indicated that patients who did better on the distinctive feature questions had more activity near the temporal pole in the left hemisphere, whether the questions were about living or nonliving things. This result is straightforwardly compatible with the concept structure

hypothesis, and is not readily explained by the localizationist position. The concept structure hypothesis can also explain why category-specific deficits for living things are observed more commonly than category-specific deficits for nonliving things. Specifically, the concept structure hypothesis suggests that the trick to discriminating living things is to pick out the few, highly specific discriminating features from among the larger number of highly correlated common features. This places greater burdens on a unified semantic processing system that handles both living and nonliving kinds (Moss and Tyler, 2003).

## Summary and Conclusions

The study of word representations and the processes that we use to activate and use stored knowledge relating to words is one of the most important enterprises in language science. Words can be represented in many different ways at many different levels of abstraction, and the way these representations are organized and connected to one another can affect the way different words are processed. An important distinction in the study of words is the difference between form and meaning. Theories of lexical access, such as logogen, TRACE, and COHORT, are primarily concerned with the representation and activation of word form information. Research in this area has shown that feature- and morpheme-level representations are an important component of the lexical access process. Although many viable models continue to be researched, and they differ in many fine details, there is general consensus that lexical access involves activation of multiple candidates and competition between them for selection. Research on lexical semantics focuses around the construct of networks of associations between words and the concepts they refer to. While we need some mechanism to tie systems of symbols to something other than symbols, associations within the lexical-symbol system appear to affect how people respond to words, and may play an important role in how new word meanings are acquired. The study of word-brain relationships indicates that word form and semantic representations are supported by widespread, left-lateralized networks; and the particular pattern of activity that occurs in response to a word depends on aspects of the word's form, its meaning, and the task that the individual is performing when the word is encountered. In terms of lexical semantics, although some categories of meaning are more vulnerable to brain damage than others, the available evidence favors the hypothesis that a single, distributed system is responsible for storing knowledge about word meanings.

### TEST YOURSELF

1. What components go together to make a word?
2. What do we mean by *meaning*? How are different meanings of the word *meaning* related?
3. How are word senses represented in long-term memory? How closely do meanings resemble dictionary definitions? How can meanings be represented in associationist networks?

4. Describe the embodied semantics hypothesis and contrast it with associationist semantics. What evidence favors each approach to lexical semantics? What role do mirror neurons play in embodied semantics?
5. Describe and contrast the logogen, FOBS, TRACE, and COHORT models of lexical access. What observations can each model account for or explain? Describe research findings that could be problematic for each account.
6. Describe the role that morphemes play in lexical access. What do priming experiments involving the manipulation of morphemes tell us about lexical access?
7. Compare the distributed cohort model to Elman's simple recurrent network model. What evidence supports each model?
8. How are ambiguous words processed?
9. What parts of the brain are involved in storing and activating information about words?
10. What is a category deficit? How do you get one? What's the best explanation for category deficit?

### THINK ABOUT IT

1. Find a newspaper article. Have a contest to see how many semantically ambiguous words you and your friends can find in the article. See if you can rewrite the first paragraph of the story so that it has no ambiguous words. Why do you think natural languages have ambiguous words? What are the advantages and disadvantages?

### Notes

- 1 Hyöna and Pollatsek (1998, p. 1612) provide the following example from Finnish of people chaining morphemes to make new words: "lumi = snow, lumipallo = snowball, lumipallosota = snowball fight, lumipallosotatantere = snowball fight field."
- 2 *Team* itself can be a token of a more general category, like *organization* (*team, company, army*). *Type* and *token* are used differently in the speech production literature. There, *token* is often used to refer to a single instance of a spoken word; *type* is used to refer to the abstract representation of the word that presumably comes into play every time an individual produces that word.
- 3 ... and *Kevin Bacon* would activate every name you know.
- 4 Which are associated because they appear together in common idioms like "It was raining cats and dogs," "They fought like cats and dogs," and because they appear together in a wide variety of scenarios.
- 5 Every word that you encounter affects your brain waves. One component of that response is an increase in negative voltage measured at the scalp that peaks about 400 ms after you hear a word. Some words cause a large peak. Some cause a smaller peak. Generally, words that make more sense in a given context produce smaller N400 effects.
- 6 Named after the homicidal computer in Stanley Kubrick's film *2001: A Space Odyssey*.
- 7 *Corpora* is the plural form of *corpus*, which is Latin for *body*.

- 8 The “foreign translator” argument is similar. As described by Glenberg and Robertson (2000, pp. 381–382), “You just landed in an airport in a foreign country and ... you do not speak the local language. As you disembark, you notice a sign printed in the foreign language (whose words are arbitrary abstract symbols to you). Your only resource is a dictionary printed in that language; that is, the dictionary consists of other arbitrary abstract symbols. You use the dictionary to look up the first word in the sign, but you don’t know the meaning of any of the words in the definition. So, you look up the first word in the definition, but you don’t know the meaning of the words in that definition, and so on. Obviously, no matter how many words you look up, that is, no matter how many structural relations you determine among the arbitrary abstract symbols, you will never figure out the meaning of any of the words.”
- 9 These kinds of ERP results need to be approached with some degree of caution because (a) the way current propagates through body tissues means that the source of the electrical activity is not necessarily directly beneath the electrode that records the response; and (b) source-modeling is an inexact science (T.Y. Swaab, personal communication).
- 10 Saygin et al. (2004) did find different patterns of reading–perception deficit correlations in different subgroups of patients, classified based on the severity of their aphasic symptoms. They concluded on the basis of perception–reading deficit correlations within mildly impaired aphasics that there may be a common neural substrate for linguistic and nonlinguistic action comprehension. They also suggested that their data support the embodied view of semantics, despite the fact that some of their strongest lesion–deficit correlations occurred for areas that are not considered to be part of the mirror neuron system.
- 11 From *logos*—“word” and *genus*—“birth” (Morton, 1969, p. 165).
- 12 But see Pinker (1994), *The Language Instinct*, for a four-letter word that gets used as an infix in colloquial English.
- 13 Obligatory cute kid story: When Rose was about 3 years old, she showed me a picture in one of her books and said, “This is an impostosaurus. It’s a dinosaur that pretends to be a different dinosaur.” The derivation is: Start with *impostor*. Strip off the *-er* suffix. Take *dinosaur*. Strip off the *-saur* suffix. Combine the two, add the *-us* suffix that goes with dinosaur species names. Pretty slick for a 3-year-old.
- 14 The revised COHORT model makes a similar claim, but only words with a semantic relationship are listed under a single lexical entry. So, *departure*, *departed*, and *departing* are stored together under the morphologically and semantically related *depart* root morpheme; but *apartment* and *apart* are stored separately, despite the fact that they both contain the common root morpheme *apart* (Marslen-Wilson et al., 1994).
- 15 Word identification in natural reading also appears to benefit from morpheme overlap between words. Word repetition and morpheme overlap had similar effects on eye movements shortly after test words were fixated. It is possible that morphological effects are stronger in languages that have richer morphological systems than English (Deutsch et al., 2003).
- 16 Foss and Swinney (1973) demonstrated that people could detect the presence of a two-syllable word *faster* than they could detect the first syllable that those words contained. So the word superiority effect extends to larger units of analysis than letters and phonemes.
- 17 One estimate suggests that 60% of English words become uniquely identifiable before their final phoneme (Luce, 1986).
- 18 In ancient Rome, large military units were divided into smaller groups of soldiers called “cohorts.”
- 19 With the exception of strongly related inflectional variants like *trespasses*, *trespassed*, *trespassing*, and so forth.
- 20 Multiple activation and limited interaction with context also occur in experimental situations where subjects are learning new vocabulary that refers to actions involving novel objects. Visual world data suggest that word onset competitors are attended to more often than other kinds of distracting stimuli, and context that restricts the possible range of referents reduces, but does not eliminate, those competition effects (Revill et al., 2008).
- 21 Logogen also does not have a straightforward means of dealing with nonword stimuli, as nonwords will not be represented by logogens (Marslen-Wilson, 1987).
- 22 Although the effect of sublexical mismatch has been claimed to be evidence against the TRACE model, more recent behavioral and modeling efforts have shown that specific versions of TRACE are capable of producing the correct pattern of coarticulatory mismatch effects (see Dahan, Magnuson, Tanenhaus et al., 2001). Such mismatch effects are still bad news for strictly bottom-up accounts like logogen.
- 23 Exception: Puns.
- 24 See Caplan (2009) for an alternative interpretation of this experiment.
- 25 Martin et al. (1996) checked for error rates by performing additional scans while subjects made an overt response, but these data could not be included in the brain activity analyses due to contamination by motor system activity involved in the speech response. Possible visual differences between tools and animals were controlled in a further experiment where subjects responded to silhouettes rather than line drawings; the brain activity results were the same.
- 26 One perplexing issue that confronts language scientists is this: Neuroimaging data suggest that frontal lobe structures play a role in the activation and use of semantic information during production and

comprehension, but many patients with frontal lobe damage do not appear to have problems performing tasks that require access to semantic information, like category judgments (e.g. Damasio et al., 1996; Price et al., 1997). One possibility is that frontal regions participate in "effortful retrieval, maintenance, and/or control of semantic information, whereas long-term storage of the conceptual and semantic knowledge is dependent on posterior regions" (Fiez, 1997, p. 81; see also Thompson-Schill et al., 1997).

- 27 There is an ongoing debate about the visual word form area. The basic question is whether the visual word form area is involved in tasks that are not related to word processing (see Schlaggar and McCandliss, 2007).

## References

- Aflalo, T., Zhang, C.Y., Rosario, E.R., et al. (2020). A shared neural substrate for action verbs and observed actions in human posterior parietal cortex. *Science Advances*, 6(43), eabb3984.
- Alloppenna, P.D., Magnuson, J.S., and Tanenhaus, M.K. (1998). Tracking the time course of spoken word recognition using eye movements: Evidence for continuous mapping models. *Journal of Memory and Language*, 38, 419–439.
- Argiris, G., Budai, R., Maieron, M., et al. (2020). Neurosurgical lesions to sensorimotor cortex do not impair action verb processing. *Scientific Reports*, 10(1), 1–15.
- Aziz-Zadeh, L., Wilson, S.M., Rizzolatti, G., et al. (2006). Congruent embodied representations for visually presented actions and linguistic phrases describing actions. *Current Biology*, 16, 1818–1823.
- Balota, D.A. and Lorch, R.F. (1986). Depth of automatic spreading activation: Mediated priming effects in pronunciation but not in lexical decision. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 12, 336–345.
- Balota, D.A., Yap, M.J., and Cortese, M.J. (2006). Visual word recognition: The journey from features to meaning (a travel update). In M.J. Traxler and M.A. Gernsbacher, *The Handbook of Psycholinguistics* (2nd ed., pp. 285–375). Elsevier.
- Barsalou, L.W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577–660.
- Bergen, B.K., Lindsay, S., Matlock, T., et al. (2007). Spatial and linguistic aspects of visual imagery in sentence comprehension. *Cognitive Science*, 31, 733–764.
- Booth, J.R., Burman, D.D., Meyer, J.R., et al. (2003). Relation between brain activation and lexical performance. *Human Brain Mapping*, 19, 155–169.
- Boulenger, V., Mechtouff, L., Thobois, S., et al. (2008). Word processing in Parkinson's Disease is impaired for action verbs but not for concrete nouns. *Neuropsychologia*, 46, 743–756.
- Boulenger, V., Roy, A.C., Paulignan, Y., et al. (2006). Cross-talk between language processes and overt motor behavior in the first 200 ms of processing. *Journal of Cognitive Neuroscience*, 18, 1607–1615.
- Bozic, M., Marslen-Wilson, W.D., Stamatakis, E.A., et al. (2007). Differentiating morphology, form, and meaning: Neural correlates of morphological complexity. *Journal of Cognitive Neuroscience*, 19, 1464–1475.
- Bradley, D. (1979). Lexical representations of derivational relations. In M. Aranoff and M. Keas (Eds.), *Juncture* (pp. 37–55). MIT Press.
- Breedin, S.D., Saffran, E.M., and Schwartz, M.F. (1998). Semantic factors in verb retrieval: An effect of complexity. *Brain & Language*, 63, 1–31.
- Bright, P., Moss, H.E., Longe, O., et al. (2006). Conceptual structure modulates anteromedial temporal involvement in processing verbally presented object properties. *Cerebral Cortex*, 16, 1066–1073.
- Buccino, G., Riggio, L., Melli, G., et al. (2005). Listening to action related sentences modulates the activity of the motor system: A combined TMS and behavioral study. *Cognitive Brain Research*, 24, 355–363.
- Buchanan, L., Westbury, C., and Burgess, C. (2001). Characterizing semantic space: Neighborhood effects in words recognition. *Psychonomic Bulletin & Review*, 8, 531–544.
- Bunn, E.M., Tyler, L.K., and Moss, H.E. (1998). Category-specific semantic deficits: The role of familiarity and property type reexamined. *Neuropsychology*, 12, 367–379.

- Burgess, C. and Lund, K. (1997). Modelling parsing constraints with high-dimensional context space. *Language and Cognitive Processes*, 12, 177–210.
- Camacho-Collados, J. and Pilehvar, M.T. (2018). From word to sense embeddings: A survey on vector representations of meaning. *Journal of Artificial Intelligence Research*, 63, 743–788.
- Caplan, D. (2009). Experimental design and interpretation of functional neuroimaging studies of cognitive processes. *Human Brain Mapping*, 30, 59–77.
- Caramazza, A. and Hillis, A.E. (1991). Lexical organization of nouns and verbs in the brain. *Nature*, 349, 788–790.
- Chao, L.L., Weisberg, J., and Martin, A. (2002). Experience-dependent modulation of category-related brain activity. *Cerebral Cortex*, 12, 545–551.
- Clarke, R. and Morton, J. (1983). Cross modality facilitation in tachistoscopic word recognition. *The Quarterly Journal of Experimental Psychology*, 35A, 76–96.
- Cleland, A.A., Gaskell, M.G., Quinlan, P.T., et al. (2006). Frequency effects in spoken and visual word recognition: Evidence from dual-task methodologies. *Journal of Experimental Psychology: Human Perception and Performance*, 32, 104–119.
- Cohen, L., Lehericy, S., Chochon, F., et al. (2002). Language-specific tuning of visual cortex? Functional properties of the visual word form area. *Brain*, 125, 1054–1069.
- Collins, A.M. and Loftus, E.F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82, 407–428.
- Collins, A.M. and Quillian, M.R. (1970). Facilitating retrieval from semantic memory: The effect of repeating part of an inference. In A.F. Sanders (Ed.), *Attention and Performance: Vol. 3* (pp. 304–314). North Holland.
- Collins, A.M. and Quillian, M.R. (1972). How to make a language user. In E. Tulving and W. Donaldson (Eds.), *Organization of Memory* (pp. 309–351). Academic Press.
- Coltheart, M., Rastle, K., Perry, C., et al. (2001). DRC: A dual route cascaded model of visual word recognition and reading aloud. *Psychological Review*, 108, 204–256.
- Connine, C.M., Blasko, D.G., and Titone, D. (1993). Do the beginnings of spoken words have a special status in auditory word recognition? *Journal of Memory and Language*, 32, 193–210.
- Corina, D.P. and Knapp, H. (2006). Sign language processing and the mirror neuron system. *Cortex*, 42, 529–539.
- Dahan, D. and Gaskell, M.G. (2007). The temporal dynamics of ambiguity resolution: Evidence from spoken-word recognition. *Journal of Memory and Language*, 57, 483–501.
- Dahan, D., Magnuson, J.S., and Tanenhaus, M.K. (2001). Time course of frequency effects in spoken word recognition: Evidence from eye movements. *Cognitive Psychology*, 42, 317–367.
- Dahan, D., Magnuson, J.S., Tanenhaus, M.K., et al. (2001). Subcategorical mismatches and the time course of lexical access: Evidence for lexical competition. *Language and Cognitive Processes*, 16, 507–534.
- Damasio, A.R. and Tranel, D. (1993). Nouns and verbs are retrieved with differently distributed neural systems. *Proceedings of the National Academy of Science*, 90, 4957–4960.
- Damasio, H., Grabowski, T.J., Tranel, D., et al. (1996). A neural basis for lexical retrieval. *Nature*, 380, 499–505.
- Davis, C.P. and Yee, E. (2021). Building semantic memory from embodied and distributional language experience. *Wiley Interdisciplinary Reviews: Cognitive Science*, 12 (5), e1555.
- Davis, M.H., Marslen-Wilson, W.D., and Gaskell, M.G. (2002). Leading up the lexical garden path: Segmentation and ambiguity in spoken word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, 28, 218–244.
- De Groot, A.M.B. (1983). The range of automatic spreading activation in word priming. *Journal of Verbal Learning and Verbal Behavior*, 22, 417–436.
- De Renzi, E. and Di Pellegrino, G. (1995). Sparing of verbs and preserved, but ineffectual reading in a patient with impaired word production. *Cortex*, 31, 619–636.
- Dehaene, S., Le Clec'H, G., Poline, J., et al. (2002). The visual word form area: A prelexical representation of visual words in the fusiform gyrus. *Neuroreport*, 13, 321–325.
- Demb, J.B., Desmond, J.E., Wagner, A.D., et al. (1995). Semantic encoding and retrieval in the left inferior prefrontal cortex: A functional MRI study of task difficulty and process specificity. *Journal of Neuroscience*, 15, 5870–5878.

- Deutsch, A., Frost, R., Pelleg, S., et al. (2003). Early morphological effects in reading: Evidence from parafoveal preview benefit in Hebrew. *Psychonomic Bulletin & Review*, 10, 415–422.
- Devlin, J.T., Russell, R.P., Davis, M.H., et al. (2002). Is there an anatomical basis for category-specificity? Semantic memory studies in PET and fMRI. *Neuropsychologia*, 40, 54–75.
- Dorfman, J. (1994). Sublexical components in implicit memory for novel words. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 1108–1125.
- Dorfman, J. (1999). Unitization of sublexical components in implicit memory for novel words. *Psychological Science*, 10, 387–392.
- Draws, E. and Zwitserlood, P. (1995). Morphological and orthographic similarity in visual word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, 21, 1098–1116.
- Dreyer, F.R. and Pulvermüller, F. (2018). Abstract semantics in the motor system?—An event-related fMRI study on passive reading of semantic word categories carrying abstract emotional and mental meaning. *Cortex*, 100, 52–70.
- Duffy, S.A., Henderson, J.M., and Morris, R.K. (1989). Semantic facilitation of lexical access during sentence processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 791–801.
- Elman, J.L. (2004). An alternative view of the mental lexicon. *Trends in Cognitive Sciences*, 8, 301–306.
- Emmorey, K.D. (1989). Auditory morphological priming in the lexicon. *Language and Cognitive Processes*, 4, 73–92.
- Erdmann, B. and Dodge, R. (1898). *Psychologische Untersuchungen über das Lesen*. M. Niemeyer.
- Farshad, M., Pavlov, Y.G., and Kotchoubey, B. (2021). Event-related potentials in an associative word pair learning paradigm. *Journal of Neurolinguistics*, 59, 101001.
- Fiez, J.A. (1997). Phonology, semantics, and the role of the left inferior prefrontal cortex. *Human Brain Mapping*, 5, 79–83.
- Forster, K.I. (1989). Basic issues in lexical processing. In W.D. Marslen-Wilson (Ed.), *Lexical Representation and Process* (pp. 75–107). MIT Press.
- Forster, K.I. and Bednall, E.S. (1976). Terminating and exhaustive search in lexical access. *Memory & Cognition*, 4, 53–61.
- Foss, D.J. and Blank, M.A. (1980). Identifying the speech codes. *Cognitive Psychology*, 12, 1–31.
- Foss, D.J. and Swinney, D.A. (1973). On the psychological reality of the phoneme: Perception, identification, and consciousness. *Journal of Verbal Learning and Verbal Behavior*, 12, 246–257.
- Friederici, A.D. (2002). Towards a neural basis of auditory sentence processing. *Trends in Cognitive Science*, 6, 78–84.
- Friederici, A.D., Opitz, B., and von Cramon, D.Y. (2000). Segregating semantic and syntactic aspects of processing in the human brain: An fMRI investigation of different word types. *Cerebral Cortex*, 10, 698–705.
- Frith, C.D., Friston, K.J., Liddle, P.F., et al. (1991). A PET study of word finding. *Neuropsychologia*, 29, 1137–1148.
- Frost, R., Forster, K.I., and Deutsch, A. (1997). What can we learn from the morphology of Hebrew? A masked-priming investigation of morphological representation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 829–856.
- Gabora, L., Rosch, E., and Aerts, D. (2008). Toward an ecological theory of concepts. *Ecological Psychology*, 20, 84–116.
- Gallese, V. and Lakoff, G. (2005). The brain's concepts: The role of the sensory-motor system in reason and language. *Cognitive Neuropsychology*, 22, 455–479.
- Gaskell, M.G. and Marslen-Wilson, W.D. (1997). Integrating form and meaning: A distributed model of speech perception. *Language and Cognitive Processes*, 12, 613–656.
- Gaskell, M.G. and Marslen-Wilson, W.D. (2001). Lexical ambiguity resolution and spoken word recognition: Bridging the gap. *Journal of Memory and Language*, 44, 325–349.
- Gaskell, M.G. and Marslen-Wilson, W.D. (2002). Representation and competition in the perception of spoken words. *Cognitive Psychology*, 45, 220–266.
- Gernsbacher, M.A. (1990). *The Structure Building Framework*. Erlbaum.

- Gernsbacher, M.A. (2009, March). *Do psycholinguists believe that humans have mirror neurons? Talk presented to the 22nd Annual CUNY Conference on Human Sentence Processing, Davis, CA.*
- Glenberg, A.M. and Kaschak, M.P. (2002). Grounding language in action. *Psychonomic Bulletin & Review*, 9, 558–565.
- Glenberg, A.M. and Robertson, D.A. (2000). Symbol grounding and meaning: A comparison of high dimensional and embodied theories of meaning. *Journal of Memory and Language*, 43, 379–401.
- Glosser, G. and Friedman, R.B. (1991). Lexical but not semantic priming in Alzheimer's disease. *Psychology and Aging*, 6, 522–527.
- Glosser, G., Grugan, P.K., Friedman, R.B., et al. (1998). Lexical, semantic, and associative priming in Alzheimer's disease. *Neuropsychology*, 12, 218–224.
- Goldberg, R.F., Perfetti, C.A., Fiez, J.A., et al. (2007). Selective retrieval of abstract semantic knowledge in left prefrontal cortex. *The Journal of Neuroscience*, 27, 3790–3798.
- Gough, P.M., Nobre, A.C., and Devlin, J.T. (2005). Dissociating linguistic processes in the left inferior frontal cortex with transcranial magnetic stimulation. *The Journal of Neuroscience*, 25, 8010–8016.
- Grabowski, T.J., Damasio, H., Tranel, D., et al. (2001). A role for left temporal pole in the retrieval of words for unique entities. *Human Brain Mapping*, 13, 199–212.
- Grosjean, F. (1980). Spoken word recognition processes and the gating paradigm. *Perception and Psychophysics*, 28, 267–283.
- Günther, F., Rinaldi, L., and Marelli, M. (2019). Vector-space models of semantic representation from a cognitive perspective: A discussion of common misconceptions. *Perspectives on Psychological Science*, 14, 1006–1033.
- Hagoort, P. (1997). Semantic priming in Broca's aphasics at a short SOA: No support for an automatic access deficit. *Brain & Language*, 56, 287–300.
- Harnad, S. (1990). The symbol grounding problem. *Physica D*, 42, 335–346.
- Hauk, O., Johnsrude, I., and Pulvermüller, F. (2004). Somatotopic representation of action words in human motor and premotor cortex. *Neuron*, 41, 301–307.
- Heim, S., Alter, K., Ischebeck, A.K., et al. (2005). The role of left Brodmann's areas 44 and 45 in reading words and pseudowords. *Cognitive Brain Research*, 25, 982–993.
- Howard, D., Patterson, K., Wise, R., et al. (1992). The cortical localization of the lexicons. *Brain*, 115, 1769–1782.
- Hutchison, K.A. (2003). Is semantic priming due to association strength or feature overlap? A microanalytic review. *Psychonomic Bulletin & Review*, 10, 785–813.
- Hyöna, J. and Pollatsek, A. (1998). Reading Finnish compound words: Eye fixations are affected by component morphemes. *Journal of Experimental Psychology: Human Perception and Performance*, 24, 1612–1627.
- Jackendoff, R. (1983). *Semantics and Cognition*. MIT Press.
- Jackson, A. and Morton, J. (1984). Facilitation of auditory word recognition. *Memory & Cognition*, 12, 568–574.
- Johnston, J.C. and McClelland, J.L. (1973). Visual factors in word perception. *Perception and Psychophysics*, 14, 365–370.
- Kiehl, K.A., Liddle, P.F., Smith, A.M., et al. (1999). Neural pathways involved in the processing of concrete and abstract words. *Human Brain Mapping*, 7, 225–233.
- Kintsch, E., Steinhart, D., Stahl, G., et al. (2000). Developing summarization skills through the use of LSA-based feedback. *Interactive Learning Environments*, 8, 87–109.
- Klatzky, R.L., Pellegrino, J.W., McCloskey, B.P., et al. (1989). Can you squeeze a tomato? The role of motor representations in semantic sensibility judgments. *Journal of Memory and Language*, 28, 56–77.
- Kluender, K.R. and Kiefte, M. (2006). Speech perception within a biologically realistic information-theoretic framework. In M.J. Traxler and M.A. Gernsbacher (Eds.), *The Handbook of Psycholinguistics* (2nd ed., pp. 153–199). Elsevier.
- Kosslyn, S.M. (1973). Scanning visual images: Some structural implications. *Perception and Psychophysics*, 14, 90–94.
- Kumar, A.A. (2021). Semantic memory: A review of methods, models, and current challenges. *Psychonomic Bulletin & Review*, 28, 40–80.

- Kuperberg, G.R., Lakshmanan, B.M., Greve, D.N., et al. (2008). Task and semantic relationship influence both the polarity and localization of hemodynamic modulation during lexico-semantic processing. *Human Brain Mapping*, 29, 544–561.
- Lahiri, A. and Marslen-Wilson, W.D. (1991). The mental representation of lexical form: A phonological approach to the recognition lexicon. *Cognition*, 38, 245–294.
- Lake, B.M. and Murphy, G.L. (2021, in press). Word meaning in minds and machines. *Psychological Review*.
- Lampe, L.F., Hameau, S., Fieder, N., et al. (2021). Effects of semantic variables on word production in aphasia. *Cortex*, 141, 363–402.
- Landauer, T.K. (1999). Latent semantic analysis: A theory of the psychology of language and mind. *Discourse Processes*, 27, 303–310.
- Landauer, T.K. and Dumais, S.T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104, 211–240.
- Landauer, T.K., Foltz, P.W., and Laham, D. (1998). An introduction to latent semantic analysis. *Discourse Processes*, 25, 259–284.
- León, J.A., Olmos, R., and Escudero, I. (2006). Assessing short summaries with human judgments procedure and latent semantic analysis in narrative and expository texts. *Behavior Research Methods*, 38, 616–627.
- Lieberman, P. (2000). *Human Language and our Reptilian Brain*. Harvard University Press.
- Lima, S.D. (1987). Morphological analysis in sentence reading. *Journal of Memory and Language*, 26, 84–99.
- Lotto, A.J., Hickok, G.S., and Holt, L.L. (2008). Reflections on mirror neurons and speech perception. *Trends in Cognitive Sciences*, 13, 110–114.
- Luce, P.A. (1986). A computational analysis of uniqueness points in auditory word recognition. *Perception and Psychophysics*, 39, 155–158.
- Lund, K. and Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instruments, & Computers*, 28, 203–208.
- Lund, K., Burgess, C., and Atchley, R.A. (1995). Semantic and associative priming in high-dimensional space. In *Proceedings of the Cognitive Science Society* (pp. 660–665). Erlbaum.
- Mahon, B.Z. and Caramazza, A. (2005). The orchestration of the sensory-motor systems: Clues from neuropsychology. *Cognitive Neuropsychology*, 22, 480–494.
- Mahon, B.Z. and Caramazza, A. (2008). A critical look at the embodied cognition hypothesis and a new proposal for grounding conceptual content. *Journal of Physiology-Paris*, 102, 59–70.
- Mainy, N., Jung, J., Baciú, M., et al. (2008). Cortical dynamics of word recognition. *Human Brain Mapping*, 29, 1215–1230.
- Marslen-Wilson, W.D. (1973). Linguistic structure and speech shadowing at very short latencies. *Nature*, 244, 522–523.
- Marslen-Wilson, W.D. (1984). Function and process in spoken word recognition. In H. Bouma and D.G. Bouwhuis (Eds.), *Attention and Performance*, vol. 10. *Control of Language Processes* (pp. 125–150). Erlbaum.
- Marslen-Wilson, W.D. (1985). Speech shadowing and speech comprehension. *Speech Communication*, 4, 55–73.
- Marslen-Wilson, W.D. (1987). Functional parallelism in spoken word recognition. *Cognition*, 25, 71–102.
- Marslen-Wilson, W.D. (1989). Access and integration: Projecting sound onto meaning. In W.D. Marslen-Wilson (Ed.), *Lexical Representation and Process* (pp. 3–24). MIT Press.
- Marslen-Wilson, W.D. (1990). Activation, competition, and frequency in lexical access. In G.T.M. Altmann (Ed.), *Cognitive Models of Speech Processing: Psycholinguistic and Computational Perspectives* (pp. 148–172). MIT Press.
- Marslen-Wilson, W.D. and Tyler, L.K. (1975). Processing structure of sentence perception. *Nature*, 257, 784–786.
- Marslen-Wilson, W.D. and Tyler, L.K. (1980). The temporal structure of spoken language understanding. *Cognition*, 8, 1–71.
- Marslen-Wilson, W.D. and Tyler, L.K. (1997). Dissociating types of mental computation. *Nature*, 387, 592–594.

- Marslen-Wilson, W.D., Tyler, L.K., Waksler, R., et al. (1994). Morphology and meaning in the English mental lexicon. *Psychological Review*, 101, 3–33.
- Marslen-Wilson, W.D. and Warren, P. (1994). Levels of perceptual representation and process in lexical access: Words, phonemes, and features. *Psychological Review*, 101, 653–675.
- Marslen-Wilson, W.D. and Welsh, A. (1978). Processing interactions and lexical access during word recognition in continuous speech. *Cognitive Psychology*, 10, 29–63.
- Marslen-Wilson, W.D. and Zwitserlood, P. (1989). Accessing spoken words: The importance of word onsets. *Journal of Experimental Psychology: Human Perception and Performance*, 15, 576–585.
- Martin, A., Wiggs, C.L., Ungerleider, L.G., et al. (1996). Neural correlates of category-specific knowledge. *Nature*, 379, 649–652.
- Mason, R.A. and Just, M.A. (2007). Lexical ambiguity in sentence comprehension. *Brain Research*, 1146, 115–127.
- McCandliss, B., Cohen, L., and Dehaene, S. (2003). The visual word form area: Expertise for reading in the fusiform gyrus. *Trends in Cognitive Sciences*, 7, 293–299.
- McClelland, J.L. and Elman, J.L. (1986). The TRACE model of speech perception. *Cognitive Psychology*, 23, 1–44.
- McClelland, J.L. and Rumelhart, D.E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review*, 88, 375–407.
- McClelland, J.L. and Rumelhart, D.E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, 114, 159–188.
- McNamara, T.P. and Altarriba, J. (1988). Depth of spreading activation revisited: Semantic mediated priming occurs in lexical decisions. *Journal of Memory and Language*, 27, 545–559.
- McQueen, J.M., Norris, D., and Cutler, A. (1994). Competition in spoken word recognition: Spotting words in other words. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 621–638.
- McQueen, J.M. and Viebahn, M.C. (2007). Tracking recognition of spoken words by tracking looks to printed words. *The Quarterly Journal of Experimental Psychology*, 60, 661–671.
- McRae, K. and Boisvert, S. (1998). Automatic semantic similarity priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, 558–572.
- McRae, K., Cree, G.S., Seidenberg, M.S., et al. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavior Research Methods*, 37, 547–559.
- McRae, K., de Sa, V.R., and Seidenberg, M.S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General*, 126, 99–130.
- Meyer, D.E. and Schvaneveldt, R.W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, 90, 227–234.
- Meyer, D.E. and Schvaneveldt, R.W. (1976). Meaning, memory structure, and mental processes. *Science*, 192, 27–33.
- Miller, J., Brookie, K., Wales, S., et al. (2018). Embodied cognition: Is activation of the motor cortex essential for understanding action verbs? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 44, 335.
- Mirman, D. and Magnuson, J.S. (2008). Attractor dynamics and semantic neighborhood density: Processing is slowed by near neighbors and speeded by distant neighbors. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34, 65–79.
- Mitchell, T.M., Shinkareva, S.V., Carlson, A., et al. (2008). Predicting human brain activity associated with the meanings of nouns. *Science*, 320, 1191–1194.
- Montero-Melis, G., Van Paridon, J., Ostarek, M., et al. (2022). No evidence for embodiment: The motor system is not needed to keep action verbs in working memory. *Cortex*, 150, 108–125.
- Morey, R.D., Kaschak, M.P., Díez-Álamo, A.M., et al. (2021). A pre-registered, multi-lab non-replication of the action-sentence compatibility effect (ACE). *Psychonomic Bulletin & Review*, 1–14.
- Morton, J. (1969). Interaction of information in word recognition. *Psychological Review*, 76, 165–178.
- Moss, H.E., Ostrin, R.K., Tyler, L.K., et al. (1995). Accessing different types of lexical semantic information: Evidence from priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 863–883.

- Moss, H.E. and Tyler, L.K. (2003). Weighing up the facts of category-specific semantic deficits. *Trends in Cognitive Sciences*, 7, 480–481.
- Moss, H.E., Tyler, L.K., Durrant-Peatfield, M., et al. (1998). “Two eyes of a see-through”: Impaired and intact semantic knowledge in a case of selective deficit for living things. *Neurocase*, 4, 291–310.
- Murphy, G.L. and Medin, D.L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289–316.
- Neely, J.H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General*, 106, 226–254.
- Negri, G.A.L., Rumiat, R.I., Zadini, A., et al. (2007). What is the role of motor simulation in action and object recognition? Evidence from apraxia. *Cognitive Neuropsychology*, 24, 795–816.
- Nelson, D.L., Bennett, D.J., Gee, N.R., et al. (1993). Implicit memory: Effects of network size and interconnectivity on cued recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19, 747–764.
- Nobre, A.C., Allison, T., and McCarthy, G. (1994). Word recognition in the human inferior temporal lobe. *Nature*, 372, 260–263.
- Nobre, A.C. and McCarthy, G. (1994). Language-related ERPs: Scalp distributions and modulation by word type and semantic priming. *Journal of Cognitive Neuroscience*, 6, 233–255.
- Nobre, A.C., Price, C.J., Turner, R., et al. (1997). Selective processing of nouns and function words in the human brain. *Neuroimage*, 5, 53.
- Noppeney, U., Patterson, K., Tyler, L.K., et al. (2007). Temporal lobe lesions and semantic impairment: A comparison of herpes simplex virus encephalitis and semantic dementia. *Brain*, 130, 1138–1147.
- Ober, B.A. and Shenaut, G.K. (2006). Semantic memory. In M.J. Traxler and M.A. Gernsbacher, *The Handbook of Psycholinguistics* (2nd ed., pp. 403–453). Elsevier.
- Ober, B.A., Shenaut, G.K., and Reed, B.R. (1995). Assessment of associative relations in Alzheimer’s disease: Evidence for the preservation of semantic memory. *Aging and Cognition*, 2, 254–267.
- Oliveri, M., Finocchiaro, A., Shapiro, K., et al. (2004). All talk and no action: A transcranial magnetic stimulation study of motor cortex activation during action word production. *Journal of Cognitive Neuroscience*, 16, 374–381.
- Onifer, W. and Swinney, D.A. (1981). Accessing lexical ambiguities during sentence comprehension: Effects of frequency of meaning and contextual bias. *Memory & Cognition*, 9, 225–236.
- Ostarek, M. and Huettig, F. (2019). Six challenges for embodiment research. *Current Directions in Psychological Science*, 28, 593–599.
- Ostrin, R.K. and Tyler, L.K. (1993). Automatic access to lexical semantics in aphasia: Evidence from semantic and associative priming. *Brain & Language*, 45, 147–159.
- Perea, M. and Gotor, A. (1997). Associative and semantic priming effects occur at very short SOAs in lexical decision and naming. *Cognition*, 67, 223–240.
- Perea, M. and Rosa, E. (2002). The effects of associative and semantic priming in the lexical decision task. *Psychological Research*, 66, 180–194.
- Petersen, S.E., Fox, P.T., Posner, M.I., et al. (1988). Positron emission tomographic studies of the cortical anatomy of single-word processing. *Nature*, 331, 585–589.
- Petersen, S.E., Fox, P.T., Posner, M.I., et al. (1989). Positron emission tomographic studies of the processing of single words. *Journal of Cognitive Neuroscience*, 1, 153–170.
- Pexman, P.M., Hargreaves, I.S., Edwards, J.D., et al. (2007). The neural consequences of semantic richness: When more comes to mind, less activation is observed. *Psychological Science*, 18, 401–406.
- Pinker, S. (1994). *The Language Instinct*. Harper.
- Posner, M.I., Petersen, S.E., Fox, P.T., et al. (1988). Localization of cognitive operations in the human brain. *Science*, 240, 1627–1631.
- Posner, M.I. and Raichle, M.E. (1994). *Images of mind*. Scientific American Books.
- Posner, M.I. and Snyder, C.R. (1975). Attention and cognitive control. In R.L. Solso (Ed.), *Information Processing and Cognition* (pp. 55–85). Erlbaum.
- Potts, G.R., Keenan, J.M., and Golding, J.M. (1988). Assessing the occurrence of elaborative inferences: Lexical decision versus naming. *Journal of Memory and Language*, 27, 399–415.

- Price, C.J., Moore, C.J., Humphreys, G.W., et al. (1997). Segregating semantic from phonological processes during reading. *Journal of Cognitive Neuroscience*, 9, 727–733.
- Pulvermüller, F. (1999). Words in the brain's language. *Behavioral and Brain Sciences*, 22, 253–336.
- Pulvermüller, F., Härle, M., and Hummel, F. (2001). Walking or talking? Behavioral and neurophysiological correlates of action verb processing. *Brain & Language*, 78, 143–168.
- Pulvermüller, F., Hauk, O., Nikulin, V.V., et al. (2005). Functional links between motor and language systems. *European Journal of Neuroscience*, 21, 793–797.
- Pulvermüller, F., Lutzenberger, W., and Preissl, H. (1999). Nouns and verbs in the intact brain: Evidence from event-related potentials and high frequency cortical response. *Cerebral Cortex*, 9, 497–506.
- Pulvermüller, F., Shtyrov, Y., and Ilmoniemi, R. (2005). Brain signatures of meaning access in action word recognition. *Journal of Cognitive Neuroscience*, 17, 884–892.
- Randall, B., Moss, H.E., Rodd, J.M., et al. (2004). Distinctiveness and correlation in conceptual structure: Behavioral and computational studies. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30, 393–406.
- Rastle, K., Davis, M.H., Marslen-Wilson, W.D., et al. (2000). Morphological and semantic effects in visual word recognition: A time-course study. *Language and Cognitive Processes*, 15, 507–537.
- Rayner, K. (1998). Eye movements in reading and information processing: Twenty years of research. *Psychological Bulletin*, 124, 372–422.
- Rayner, K. and Duffy, S.A. (1987). Eye movements and lexical ambiguity. In J.K. O'Regan and A. Levy-Schoen, *Eye Movements: From Physiology to Cognition* (pp. 521–529). Elsevier.
- Rayner, K. and Pollatsek, A. (1989). *The Psychology of Reading*. Erlbaum.
- Rayner, K. and Pollatsek, A. (2006). Eye-movement control in reading. In M.J. Traxler and M.A. Gernsbacher (Eds.), *The Handbook of Psycholinguistics* (2nd ed., pp. 613–657). Elsevier.
- Reicher, G.M. (1969). Perceptual recognition as a function of meaningfulness of stimulus material. *Journal of Experimental Psychology*, 81, 274–280.
- Revill, K.P., Tanenhaus, M.K., and Aslin, R.N. (2008). Context and spoken word recognition in a novel lexicon. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34, 1207–1223.
- Rhodes, S.M. and Donaldson, D.I. (2008). Association and not semantic relationships elicit the N400 effect: Electrophysiological evidence from an explicit language comprehension task. *Psychophysiology*, 45, 50–59.
- Rips, L.J., Shoben, E.J., and Smith, E.E. (1973). Semantic distance and the verification of semantic relations. *Journal of Verbal Learning & Verbal Behavior*, 12, 1–20.
- Rissman, J., Eliassen, J.C., and Blumstein, S.E. (2003). An event-related fMRI investigation of implicit semantic priming. *Journal of Cognitive Neuroscience*, 15, 1160–1175.
- Rizzolatti, G. and Arbib, M.A. (1998). Language within our grasp. *Trends in Neurosciences*, 21, 188–194.
- Rizzolatti, G. and Craighero, L. (2004). The mirror-neuron system. *Annual Review of Neuroscience*, 27, 169–192.
- Rodd, J.M., Gaskell, M.G., and Marslen-Wilson, W.D. (2004). Modelling the effects of semantic ambiguity in word recognition. *Cognitive Science*, 28, 89–104.
- Rosch, E.H. (1973). Natural categories. *Cognitive Psychology*, 4, 328–350.
- Roskies, A.L., Fiez, J.A., Balota, D.A., et al. (2001). Task-dependent modulation of regions in the left inferior frontal cortex during semantic processing. *Journal of Cognitive Neuroscience*, 13, 829–843.
- Rumelhart, D.E. and McClelland, J.L. (1982). An interactive activation model of context effects in letter perception: Part 2. The contextual enhancement effect and some tests and extensions of the model. *Psychological Review*, 89, 60–94.
- Saffran, E.M. and Schwartz, M.F. (1994). Of cabbages and things: Semantic memory from a neuropsychological perspective: A tutorial review. In C. Umiltà and M. Moscovitch (Eds.), *Attention and Performance: Vol. 15. Conscious and Nonconscious Information Processing* (pp. 507–536). MIT Press.
- Salverda, A.P., Dahan, D., and McQueen, J.M. (2003). The role of prosodic boundaries in the resolution of lexical embedding in speech comprehension. *Cognition*, 90, 51–89.

- Salverda, A.P., Dahan, D., Tanenhaus, M.K., et al. (2007). Effects of prosodically modulated sub-phonetic variation on lexical competition. *Cognition*, 105, 466–476.
- Savin, H.B. and Bever, T.G. (1970). The nonperceptual reality of the phoneme. *Journal of Verbal Learning and Verbal Behavior*, 9, 295–302.
- Saygin, A.P., Wilson, S.M., Dronkers, N.F., et al. (2004). Action comprehension in aphasia: Linguistic and non-linguistic deficits and their lesion correlates. *Neuropsychologia*, 42, 1788–1804.
- Schlaggar, B.L. and McCandliss, B.D. (2007). Development of neural systems for reading. *Annual Review of Neuroscience*, 30, 475–503.
- Scott, S.K., Blank, C.C., Rosen, S., et al. (2000). Identification of a pathway for intelligible speech in the left temporal lobe. *Brain*, 123, 2400–2406.
- Searle, J.R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3, 417–457.
- Seidenberg, M.S. and Tanenhaus, M.K. (1979). Orthographic effects on rhyme monitoring. *Journal of Experimental Psychology: Human Learning and Memory*, 5, 546–554.
- Seidenberg, M.S., Tanenhaus, M.K., Leiman, J.M., et al. (1982). Automatic access of the meanings of ambiguous words: Evidence from priming and eye fixations. *Memory & Cognition*, 28, 1098–1108.
- Setola, P. and Reilly, R.G. (2005). Words in the brain's language: An experimental investigation. *Brain & Language*, 94, 251–259.
- Shalom, D.B. and Poeppel, D. (2008). Functional anatomic models of language: Assembling the pieces. *The Neuroscientist*, 14, 119–127.
- Shatzman, K.B. and McQueen, J.M. (2006). Prosodic knowledge affects the recognition of newly acquired words. *Psychological Science*, 17, 372–377.
- Shelton, J.R. and Martin, R.C. (1992). How semantic is automatic semantic priming? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 1191–1210.
- Skinner, B.F. (1957). *Verbal Behavior*. Copley Publishing Group.
- Smith, E.E., Shoben, E.J., and Rips, L.J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review*, 81, 214–241.
- Smith, P.T. and Sterling, C.M. (1982). Factors affecting the perceived morphological structure of written words. *Journal of Verbal Learning and Verbal Behavior*, 21, 704–721.
- Speed, L.J. and Majid, A. (2020). Grounding language in the neglected senses of touch, taste, and smell. *Cognitive Neuropsychology*, 37, 363–392.
- Stanners, R.F., Neiser, J.J., and Painton, S. (1979). Memory representation for prefixed words. *Journal of Verbal Learning & Verbal Behavior*, 18, 733–743.
- Stevens, A. and Coupe, P. (1978). Distortions in judged spatial relations. *Cognitive Psychology*, 10, 422–437.
- Swinney, D.A. (1979). Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*, 18, 645–659.
- Tabossi, P., Colombo, L., and Job, R. (1987). Accessing lexical ambiguity: Effects of context and dominance. *Psychological Research*, 49, 161–167.
- Tabossi, P. and Zardon, F. (1993). Processing ambiguous words in context. *Journal of Memory and Language*, 32, 359–372.
- Taft, M. (1979). Recognition of affixed words and the word frequency effect. *Memory & Cognition*, 7, 263–272.
- Taft, M. (1981). Prefix stripping revisited. *Journal of Verbal Learning and Verbal Behavior*, 20, 289–297.
- Taft, M. (1994). Interactive activation as a framework for understanding morphological processing. *Language and Cognitive Processes*, 9, 271–294.
- Taft, M. and Forster, K. (1975). Lexical storage and retrieval for prefixed words. *Journal of Verbal Learning and Verbal Behavior*, 14, 638–647.
- Tettamanti, M., Buccino, G., Saccuman, M.C., et al. (2005). Listening to action-related sentences activates fronto-parietal motor circuits. *Journal of Cognitive Neuroscience*, 17, 273–281.
- Thompson-Schill, S.L., Aguirre, G.K., D'Esposito, M., et al. (1999). A neural basis for category and modality specific semantic knowledge. *Neuropsychologia*, 37, 671–676.

- Thompson-Schill, S.L., D'Esposito, M., Aguirre, G.K., et al. (1997). Role of left inferior prefrontal cortex in retrieval of semantic knowledge: A reevaluation. *Proceedings of the National Academy of Science*, 94, 14792–14797.
- Tomasino, B., Fink, G.R., Sparing, R., et al. (2008). Action verbs and the primary motor cortex: A comparative TMS study of silent reading, frequency judgments, and motor imagery. *Neuropsychologia*, 46, 1915–1926.
- Tranel, D., Grabowski, T.J., Lyon, J., et al. (2005). Naming the same entities from visual or from auditory stimulation engages similar regions of left inferotemporal cortices. *Journal of Cognitive Neuroscience*, 17, 1293–1305.
- Tranel, D., Logan, C.G., Frank, R.J., et al. (1997). Explaining category-related effects in the retrieval of conceptual and lexical knowledge for concrete entities: Operationalization and analysis of factors. *Neuropsychologia*, 35, 1329–1339.
- Tucker, M. and Ellis, R. (2001). The potentiation of grasp types during visual object categorization. *Visual Cognition*, 8, 769–800.
- Tucker, M. and Ellis, R. (2004). Action priming by briefly presented objects. *Acta Psychologica*, 116, 185–203.
- Tyler, L.K. and Wessels, J. (1983). Quantifying contextual contributions to word-recognition processes. *Perception and Psychophysics*, 34, 409–420.
- Vainio, L., Symes, E., Ellis, R., et al. (2008). On the relations between action planning, object identification, and motor representations of observed actions and objects. *Cognition*, 108, 444–465.
- Vandenberghe, R., Price, C., Wise, R., et al. (1996). Functional anatomy of a common semantic system for words and pictures. *Nature*, 383, 254–256.
- Vannuscorps, G. and Caramazza, A. (2019). Conceptual processing of action verbs with and without motor representations. *Cognitive Neuropsychology*, 36, 301–312.
- Vitale, F., Padrón, I., Avenanti, A., et al. (2021). Enhancing motor brain activity improves memory for action language: A tDCS study. *Cerebral Cortex*, 31, 1569–1581.
- Wagner, A.D., Desmond, J.E., Domb, J.B., et al. (1997). Semantic repetition priming for verbal and pictorial knowledge: A functional MRI study of left inferior prefrontal cortex. *Journal of Cognitive Neuroscience*, 9, 714–726.
- Wheeler, D.D. (1970). Processes in word recognition. *Cognitive Psychology*, 1, 59–85.
- Wible, C.G., Han, S.D., Spencer, M.H., et al. (2006). Connectivity among semantic associates: An fMRI study of semantic priming. *Brain & Language*, 97, 294–305.
- Zipf, G.K. (1949). *Human Behavior and the Principle of Least Effort*. Hafner.
- Zwaan, R.A., Madden, C.J., Yaxley, R.H., et al. (2004). Moving words: Dynamic representations in language comprehension. *Cognitive Science*, 28, 611–619.
- Zwaan, R.A. and Rapp, D.N. (2006). Discourse comprehension. In M.J. Traxler and M.A. Gernsbacher (Eds.), *The Handbook of Psycholinguistics* (2nd ed., pp. 725–764). Elsevier.
- Zwaan, R.A. and Taylor, L.J. (2006). Seeing, acting, understanding: Motor resonance in language comprehension. *Journal of Experimental Psychology: General*, 135, 1–11.
- Zwitserslood, P. (1989). The locus of the effects of sentential-semantic context in spoken-word processing. *Cognition*, 32, 25–64.