7 Word Recognition

In Chapter 5, we explored the ways in which babies and young children learn to pluck word chunks out of the streams of speech that they hear, pair these chunks up with meanings, and add to a growing collection of words that will eventually tally up to tens of thousands, if not more. But learning words and their meanings is only part of the story. Once we've added these thousands of words to our memory's stash of vocabulary units, we also need to be able to quickly and efficiently retrieve them while producing and understanding language.

Both speaking and comprehending language involve matching up meanings with their corresponding sound sequences at breakneck speed while at the same time juggling other complicated information. During a normal conversation, speech unfurls at a rate of about 150 words per minute; master debaters and auctioneers can reach speeds of 400–500 words per minute. This means that while listening to your conversational partner speaking at a typical leisurely rate, you have less than half a second to match up the sounds coming out of her mouth and rummage around in your voluminous vocabulary to find exactly the right word before the next one follows hot on its heels, all while working out how that word figures in the syntactic structure and meaning of the unfolding sentence. Most of the time you're doing all this while also planning your own response or interjection.

Perhaps it's not entirely surprising that, as a species that depends on transmitting information via language, we've evolved to have an efficient set of routines that allow us to pack and unpack large volumes of spoken information in very small spans of time. Presumably, we've had many generations—likely more than 100,000 years' worth—over which to develop the brains for it, and a knack for speedy language comprehension would seem to offer certain evolutionary advantages. On the other hand, *written* language is a fairly recent human innovation, with the oldest known writing system dating back a mere 5,000 years, to about 3200 B.C. Even today, many (hearing) societies in the world get by without a written language, though no society is speechless. Nevertheless, our ability to strip meaning from symbols on a page is at least as fast as our ability to do so with spoken sounds. The pleasure reading of skilled readers proceeds at a clip of about 200–400 words per minute—roughly 3–6 words per *second*.

The recognition of spoken words and the reading of written language offer quite different challenges from the standpoint of information processing. In speech, sounds have to be correctly identified, word boundaries have to be accurately located, and a string of sounds has to be mapped onto the correct word, despite the fact that many other words might offer *almost* as good a match as the right one. In writing, a jumble of arbitrary symbols has to be mapped onto words that have usually been learned in terms of their sound, not in terms of their visual symbols. The system of visual symbols represents a whole new set of symbolic units and mappings that's been artificially grafted onto the "natural" system of spoken language. Skilled reading relies on smoothly integrating this artificial system with the "natural" one during normal language processing—perhaps it's a little bit like the linguistic equivalent of learning to use a prosthetic limb.

In the previous paragraphs, I've been a bit preoccupied with the *time* it takes to recognize words. As it happens, time is the central obsession of the researcher who studies word recognition or, more generally, language processing. This preoccupation with time goes far beyond mere trivia, or even the desire to explain how it is that we process language as quickly as we do. Specific theories of language processing stand or fall based on the predictions they make about the relative timing of certain processing events. Time also serves as one of the most important methodological tools for studying how language processing works. As you'll see, researchers have come to learn a great deal about *how* linguistic information is processed by looking in minute detail at *how long* people take to process it, and by making careful comparisons of the time it takes to process stimuli that differ along specific dimensions. We owe much of what we know about word recognition to clever and meticulous forms of timekeeping. In this chapter, you can get a small taste of what it's like to do research on word recognition by taking part in the many Web Activities that are sprinkled throughout. A number of these activities will focus your attention on the precision and attention to detail that's needed in order to construct experiments where the results often hinge on finding differences of less than a tenth of a second of processing time.

7.1 A Connected Lexicon

Word webs

What happens in your mind when you hear or read a word? In Chapter 1, I emphasized the fact that our subjective intuitions about language often miss some critical information, and as you'll see, that's certainly the case when it comes to understanding something as basic as word recognition. Let's start by trying to describe the very simplest case of recognizing a word in isolation, outside of the context of a sentence. Based on how the experience *feels*, we might describe the process as a bit like this: Retrieving words from memory is like getting words out of a vending machine—let's think of word representations as the snacks you're trying to buy. Specific sequences of sounds you hear or letters you read are like the sequences of letters and numbers you have to punch into the

machine to get the right product to come out. Just as a vending machine has a program set up to link a sequence of button presses with a specific location that delivers your chosen snack to you, letters or sounds are programmed in your mind to activate a specific word representation. In both cases, what you want conveniently drops down for your use in the form of a single, packaged unit. When it's described this way, there doesn't seem to be much to the process of retrieving words; it's simply a matter of punching in the right input in terms of sounds or letters, which is then linked in a one-to-one fashion with the right word representation. Why *shouldn't* this happen very quickly in real time?

The truth is that this subjective impression of word recognition is deeply wrong. In actual fact, recognition of either spoken or written words is quite a bit messier than this. It seems that words aren't organized in our minds independently of one another, but rather, are connected together in complex webs. When we try to retrieve one word, we end up pulling a string that has the actual matching word, but also has a bunch of connected words dangling from it as well. But since we have the *impression* that we've pulled a single word out of our mind, what is it that has led language researchers to the conclusion that words are actually highly interconnected? To get there, they've had to find ways to probe aspects of the word recognition process that may not be completely accessible to conscious intuition.

Evidence for partial retrieval of related words

One of the strongest (and earliest) sources of evidence for the interconnections among words has come from the phenomenon of **semantic priming**, which suggests that when you hear or read a word, you also partially activate other words that are related in meaning. The relevant experiments (the first of which were performed by David Meyer and Roger Schvaneveldt, 1971) typically use a method known as the **lexical decision task**, in which participants read strings of letters on a screen that are either actual words (for example, *doctor*) or nonsense words (*domter*) and then press one of two buttons—one button if they think they've seen a real word, and another to signal that the letters formed a nonsense word. The speed with which subjects press each button is recorded.

Semantic priming experiments have shown that participants are faster to recognize a real word if it follows hard on the heels of a word that's related in meaning. For example, responses to the test word *doctor* would be speedier if that word occurred in the experiment just after the word *nurse* than if it followed an unrelated word like *butter*. This suggests that the word *nurse* wasn't accessed in isolation. If you think of word recognition as "lighting up" a word in your mind, then some of the light spills over onto neighboring words in the semantic space. This partial lighting up of *doctor* following *nurse* makes it easier to recognize *doctor* when you see that word on the screen.

The lexical decision task shows that this happens with written words. But spoken language also seems to activate related words in much the same way. One way to see quite vividly what's happening as people process spoken words is to track their eye movements to a visual scene in real time as people hear instructions related to that scene. For example, imagine seeing a screen showing pictures of, say, a hammer, a nail, a cricket, and a box of tissues, and then hearing an instruction to "click on the hammer." Seems like a trivial enough task. But a record of subjects' eye movements (see **Figure 7.1**; Yee & Sedivy, 2006) shows that as they're hearing *hammer*,

semantic priming The phenomenon by which hearing or reading a word partially activates other words that are related in meaning to that word, making the related words easier to recognize in subsequent encounters.

lexical decision task An experimental task in which participants read strings of letters on a screen that might either be actual words (*doctor*) or nonsense words (*domter*). Subjects press one button if they think they've seen a real word, or a different button to signal that the letters formed a nonsense word. Response times for real words are taken as a general measure of the ease of recognizing those words under specific experimental conditions.

WEB ACTIVITY 7.1

Interconnected words In this activity, you'll explore how to ascertain which words might be linked together in memory, through a pen-and-paper word association test in which people are asked to list the first words that come to mind when they encounter a word.

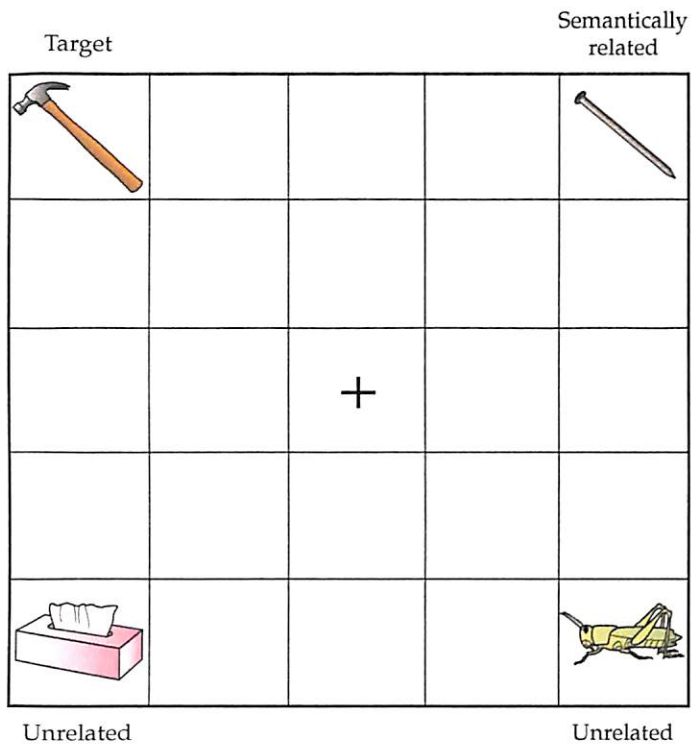

Figure 7.1 A sample display from an experiment by Yee and Sedivy (2006). Shortly after hearing the word *hammer*, subjects were more likely to look briefly at the nail than at unrelated items such as the cricket. This suggests that hearing a word results in the activation of semantically related words as well as the target word.

they might be briefly lured by a picture of a related word—that is, they're more likely to look at the nail than at the unrelated cricket. This aligns nicely with the evidence from lexical decision tasks, suggesting that words that are semantically related to the target word get an extra burst of activation, relative to other words in the lexicon.

Competition from partially activated words

The lexical decision experiments show evidence of **facilitation**: a word is made easier to recognize by having it occur after a semantically related one. But as you might imagine, if the process of retrieving a word partly lights up several other words, this could sometimes make it more *difficult* to pick out the correct word from among the other, partially lit-up ones. That is, related words that are partially activated could get in the way of recognizing the right one, a result that is referred to as **inhibition**. The experimental record shows evidence of this too, especially when it comes to words that are similar in *form*, rather than meaning. For example, if "primed" with the word *stiff*, people might be slower to subsequently recognize *still* as a word (Slowiaczek & Hamburger, 1992). In this case, the prime word seems to *compete* with the target word, tripping up the recognition process. Again, the data from the priming paradigm is corroborated by eye-tracking experiments; if the visual array includes a beetle as well as a beaker, for instance, recognition of the spoken word *beaker* tends to be slowed down, so it takes people longer to locate the image of the beaker in the display (Allopenna et al., 1998).

Competition among lexical items in memory can be seen easily when the form of a target word is similar to other words that are relevant in the immediate context; for example, in the lexical decision priming paradigm, a word might inhibit the access of a similar one immediately following it, or in the eye-tracking paradigm, locating the target referent for a word may be slowed down because the eye is drawn to an object corresponding to a similar-sounding word. But these are somewhat contrived experimental scenarios in which similar words have been deliberately planted. Most of the time when we're retrieving words from memory in real-life conversation, we don't have to cope with such blatant decoys—for example, how often would we be in a context where both *beetle* and *beaker* would be uttered? So, to what extent is lexical competition likely to really play a role in our routine language comprehension?

Quite a bit, it would seem. Competition among words happens even when there are no similar words in the immediate context—it's enough merely for there to be similar words in your own personal mental storehouse. The evidence comes from comparing how quickly people retrieve words that have either many or few sound-alikes in the general lexicon, even when those sound-alike words aren't prominent in the context. As it happens, some words that you know are relatively unique when it comes to their sound structure. For instance, take the word *stench*: try to come up with as many other words as you can that differ from that one by only one sound. After *staunch* and the more uncommon word *stanch* (as in: *stanch the flow of blood*), I pretty much draw a blank myself. But I can reel off quite a few examples that differ by only one sound from the word *sling*: *sting, fling, bling, cling, slung, slang, slim* (remember, *-ng* counts as one *sound*), *slit, slip, slid, slick*. Psycholinguists have invoked a real-estate metaphor to describe this difference between *stench* and *sling*; we say that *stench* is found in a "sparse neighborhood," with very few sound-based neighbors, while *sling* resides in a "dense neighborhood," with its neighbors crowding up on all sides.

facilitation Processes that make it easier for word recognition to be completed.

inhibition Processes that result in word recognition becoming more difficult.

WEB ACTIVITY 7.2

Neighborhood density In this activity, you'll generate words that belong to dense and sparse neighborhoods, and you'll consider other variables that threaten to mask the effect of neighborhood density if not properly controlled.

Again, psycholinguists have used response times in lexical decision tasks to probe for **neighborhood density effects**. All things being equal, after the length and frequency of words has been controlled for (since these factors can also affect how long it takes to recognize a word), people take longer to recognize words that come from dense neighborhoods than words that come from sparse neighborhoods (e.g., Goldinger et al., 1989). This effect can be seen even though no similar words have occurred before the test word, suggesting that lexical neighbors can compete for recognition even when they are not present in the immediate context. It's enough that they're simply part of the vocabulary.

neighborhood density effects Experimental results demonstrating that it is more difficult and time-consuming to retrieve a word from memory if the word bears a strong phonological resemblance to many other words in the vocabulary than if resembles only a few other words.

Building a model of word recognition

Vague metaphors involving spotlights or neighborhoods are useful for getting an intuitive sense of how words in the mental lexicon affect each other in the process of word recognition. But it's also worth moving toward more precise theories of what lies beneath these effects. Making explicit models can advance a field very quickly, whether they're communicated as simple diagrams or actually implemented as working computer programs that simulate real language processes. A model not only serves as a way of explaining current experimental results, but also tends to force researchers to grapple with other perfectly reasonable ways to account for the same data. When you have to make a decision about the details of your model, you become aware of the things you don't yet know. So the process of building models is really useful for throwing light onto as-yet-unanswered questions that researchers might otherwise not think of. One of the main benefits of a model is that it makes new predictions that can be tested experimentally.

There's probably no other psycholinguistic process that's been modeled in as much detail and by so many different rivals as word recognition. These modeling efforts have led to an enormous volume of published experimental papers dealing with word recognition. There isn't room in this chapter to give a thorough overview of all of the models that have been proposed, or the experimental evidence that has been generated. Instead, my goal will be to bring out some of the key findings and big questions about how words are processed in the mind, and to relate these to the choices that go into building a good working model.

Let's start with one simple way to model how related words affect each other in the process of word recognition, as sketched out in Figure 7.2. Here, rather than just being listed as isolated units, words are represented as belonging to a complex network. Words that are related in meaning are linked together in the network—for example, *lion* is linked to *mane* and *tiger*, and *tiger* is linked to *stripes*. (Links could also be established between words that frequently occur together in speech, whether or not their meanings are similar—for example, between *birthday* and *cake*, as illustrated by the word-association task in Web Activity 7.1). If a match is made between the sounds (or letters) being perceived and a particular word, that word gets a surge of ener-

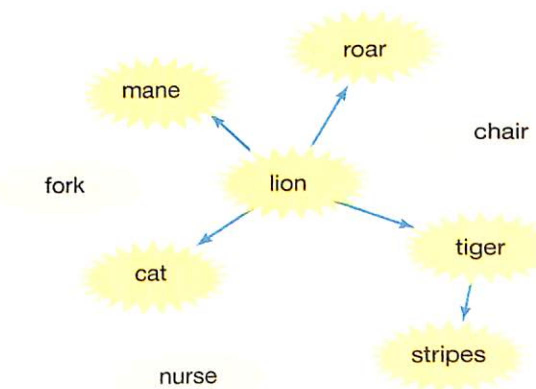

Figure 7.2 A simple spreading-activation network. Hearing the word *lion* activates the mental representation of that word and, in turn, the representations of related words such as *mane*, *cat*, *tiger*, and eventually *stripes* (via *tiger*).

mediated semantic priming The process by which a prime word (e.g., *lion*) speeds up responses to a target word (e.g., *stripes*) not because of a direct connection between *lion* and *stripes*, but due to an indirect connection via some other intervening word (e.g., *tiger*).

gy or activation. Since the words in the lexicon aren't isolated from one another but share links with quite a few other words, activation "flows" throughout the network along these links. For example, hearing or reading the word *lion* will also excite the connected words *mane* and *tiger*. This has behavioral consequences, such as the priming effects we've just seen.

Setting words up in a network like this (we've just created a mini-model!) results in a new prediction: Since each word is connected to a number of other words, and energy spreads along links in the network, we should see signs of activation traveling along paths of interconnected words. That is, *lion* should activate not just *tiger*, but also the word *stripes*. The latter word isn't related to *lion*, but it *is* related to *tiger*, which is itself linked to the perceived word. In fact, there is evidence for just this kind of **mediated semantic priming**, in which a prime word like *lion* speeds up responses to a target word like *stripes*, even though the prime and target are only related via some other intervening word.

Our simple spreading-activation story looks promising, but it's going to need some adjustments and refinements. So far, there's nothing to prevent activation from continuing to spread throughout the network from link to link to link in such a way that it eventually activates just about every word in the network. For example: *lion* → *tiger* → *stripes* → *paisley* → *shirt* → *tie* → *neck* → *head* → *hair*, and so on. The model needs to have a way to prevent such overwhelming buzzing within the lexical network.

One way to achieve this is by building in a *decay function* by which activation levels gradually die down over time. This means that the activation of *lion* would surge at first, but then fade. Since it would take some time for activation to spread throughout the network, the decay function would limit how much energy is passed on to remotely connected words or concepts. So, the activation of *lion* would have dwindled somewhat by the time it spread to *tiger*, meaning that *tiger* would be less activated than *lion* was initially. As a result, *stripes* would receive less activation than *tiger*, and any activation that *stripes* passed on to *paisley* might well be negligible by that point. With this modification, it should now be possible to quantitatively tinker with the specific rates at which activation spreads and decays so that the model will closely simulate the patterns of results from experiments with humans. For example, the model should now also be able to capture the fact that the degree to which a related prime speeds up responses to a target word depends not only on how closely related the two words are, but also on how much time has elapsed between the presentation of the prime and target words.

So far, we have a model that does a nice job of capturing *facilitatory* priming effects and their limits. But what about the cases in which the presence of related words *slows down* or impedes the recognition of the target? As it stands, the model doesn't predict these competition effects, so we need to either scrap it or add something else. Given that we already have some mileage from our little model in accounting for priming effects, the latter strategy seems like a good place to start.

Competition effects are seen most clearly among words that are related to each other in form rather than meaning, so we might focus on the process of mapping sounds or letters to word representations. This is depicted in Figure 7.3A, where phonemic (or orthographic) units are connected to word representations. When a sound (or letter) is identified, it becomes activated, and by virtue of the connections it has to words, activation flows to those words that contain it. So far, this predicts that words that contain overlapping sounds (for example, *beam* and *beat*) should both become active. But there's nothing in the model yet to explain why hearing *beam* should make it harder to subsequently recognize *beat*—in fact, as it stands, the model predicts the opposite,

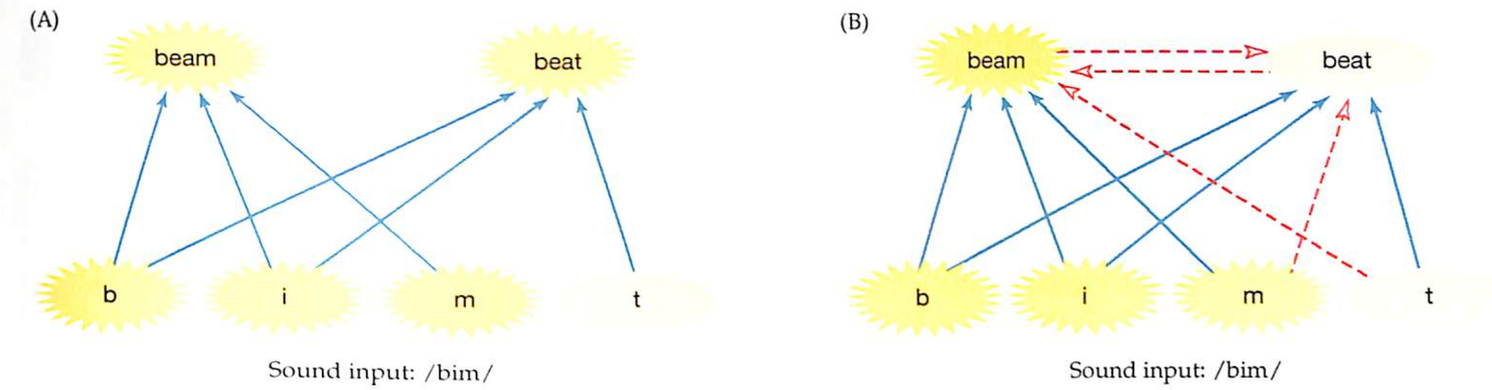

Figure 7.3 (A) A simple model showing only excitatory connections (solid lines and arrows) from the phonemic units in the words *beam* and *beat* (/bim/ and /bit/ respectively). With only excitatory connections, there's nothing to explain why hearing the word *beam* should make it harder to subsequently recognize *beat*. (B) A model with inhibitory connections (red dashed lines) in addition to excitatory links (solid lines and arrows). As the activation of a phonemic or word unit rises, activation is decreased for units that are connected to it via inhibitory links. For example, the rise in activation for *beam* results in the suppression of activation of *beat*.

since connections allow activation to spread. We can get the right result if we propose that two kinds of connections can exist between representations that share a link: **excitatory connections**, which pass activation from one unit to another, and **inhibitory connections**, which have exactly the opposite effect so that the more active a unit is, the more it *suppresses* the activation of a unit it is linked to. You can see this in the revised model in Figure 7.3B; in this version of the model, once the last sound or letter of *beam* is perceived, it will excite the word *beam*. But *beat* will be inhibited: as the activation of the unit *m* rises, it will suppress the word *beat* through the inhibitory link between *m* and *beat*; and *beat* will also become inhibited by virtue of the inhibitory link from the word *beam* to *beat*.

excitatory connections Connections along which activation is passed from one unit to another, so that the more active a unit becomes, the more it increases the activation of a unit it is linked to.

inhibitory connections Connections that lower the activation of connected units, so that the more active a unit becomes, the more it *suppresses* the activation of a unit it is linked to.

As you've seen, a great deal of the evidence for the degree of activation of word representations comes from either eye movement studies or experiments using the lexical decision task. In many cases, researchers strive to test predictions that involve very subtle differences in activation between types of stimuli, or changes in activation levels for the same stimuli over time. Hence, the experimental methods that they use need to be deployed with great precision. Method 7.1 lays out some of the challenges that come up in using the lexical decision task, and the techniques that researchers use to get the most out of this simple task.

WEB ACTIVITY 7.3

Tinkering with models and making predictions In this activity, you'll try your hand at making empirical predictions based on subtle variations in the details of the model you've just seen.

Probing the model's assumptions

Our model-building exercise so far has focused on two aspects of the model: (1) the existence of links that connect representations to one another, and (2) what happens along these links (that is, do they spread activation or suppress it, and how do these effects dissipate over time?). But our little model also makes certain implicit commitments that we haven't yet defended empirically.

We've assumed that the input to word recognition is a set of discrete letter or sound units, represented as integral units rather than as bundles that

METHOD 7.1

Using the lexical decision task

There's a simple logic behind the semantic priming techniques that are used in lexical decision tasks. You create an experiment in which you record and compare the response times to test words, or targets, that have been preceded by primes: words that are either *related* or *unrelated* to the target. For example, response times to the target *doctor* preceded by the prime *nurse* would be compared with results for *doctor* preceded by *chair*. If responses to the former are significantly faster than to the latter, you've found a priming effect, suggesting that seeing the word *nurse* facilitated the subject's recognition of the word *doctor* through spreading activation. But actually running an experiment using priming techniques involves making many small technical decisions, each of which could conceivably have some impact on your results or how you interpret them.

First, you'll need to deal with the possibility that participants might try to anticipate patterns within the experiment in order to respond strategically. If this happens, their response times might say less about what people typically do in daily conversation than about what students do when trying to puzzle out a psychology experiment. The goal is to minimize patterns wherever possible. For example, you can easily eliminate the expectation that the correct response to the target will be to press the button for "Yes, it's a word" by making sure that you put in plenty of fillers, and balance the experiment so that the likelihood of the target being a real word is exactly 50% over the course of the experiment.

But some patterns are impossible to eliminate outright, such as the relationship that exists between some of the primes and targets. If, after a while, your participants begin to catch on to the fact that the target is often related to the prime, they might approach the task less as a word recognition task, and more as a word association task: once they see the prime *nurse*, for example, they might start actively thinking of related words. In that case, it would be unsurprising to find that there's a close relationship between word association tests and priming patterns (see Web Activity 7.1); but you'd be less confident in concluding that related words are *spontaneously* activated during routine word recognition out in the wild.

You can rely on two approaches to reduce the possibility that participants will consciously notice the relationship between primes and targets. The first is to simply include a great number of filler items in which the prime and target aren't related in any way, making it harder to detect the pattern. For example, it would be harder to notice the relationships between words in this set of stimuli:

FREEDOM—METAL
WRENCH—BOOK
HANDLE—SHOES
NURSE—DOCTOR
FLOWER—SCREEN
PAPER—ROOF

than in this one:

FREEDOM—METAL
WRENCH—HAMMER
HANDLE—DOOR
NURSE—DOCTOR
FLOWER—VASE
PAPER—ROOF

METHOD 7.1 *(continued)*

The other way to reduce strategic responding is to provide your participants with as little time as possible to anticipate words that are related to the prime. You can do this by shrinking the time between the presentation of the prime and the presentation of the target. In fact, you can even exploit an intriguing quirk of the human perceptual system, namely, the fact that there's a time lag between conscious and unconscious perceptual processing, so at very rapid speeds, it's possible for our unconscious minds to have processed a stimulus while our conscious minds don't "know" that we've done so.

Researchers have used subliminal presentation of the prime word—in a paradigm known as **masked priming**—as a tool to eliminate the possibility of conscious, strategic responding. At the beginning of a trial, participants see a row of dashes or # symbols "masking" the prime. The prime is then briefly flashed and quickly covered up again before the presentation of the target stimulus. For example:

mask	#####	(1,000 milliseconds)
prime	NURSE	(50 ms)
mask	#####	(500 ms)
target	DOCTOR	(response required—is it a word?)

As we've seen from our discussion of modeling, the interconnected lexicon is in a constant state of flux, reflecting the spreading and dampening of activation over time along various links. This means that using the priming technique—which probes for responses at *one* specific time point—is a bit like relying on a snapshot to capture an object that is in motion. Depending on when you press the shutter, you may get a very different picture of the object's path of motion. For instance, if you're interested in finding out whether target words that have some particular relationship to their primes become activated above baseline levels (that is, relative to targets with unrelated primes), whether or not you find a priming effect may depend upon when you present the target. Probe too soon, and activation may not have spread to the target yet. Probe too late, and activation may have dissipated or become suppressed.

For this reason, researchers often sample at different time points, by varying either how long the prime is presented or the **interstimulus interval (ISI)**, which refers to the amount of time between the offset of the prime and the onset of the target.

Another issue with using a snapshot technique to capture a temporally dynamic process is that response times will be affected by how deeply people process the target words before making a decision. If they press the button after very shallow processing, they may not yet have had time to fully access the word's semantic representation. This can happen if the non-word targets are very distinct from any possible words (for example, *bgltx*, *aoitvb*), so decisions about the target's status can be made very quickly and on the basis of very superficial characteristics. On the other hand, if the non-words look like possible though non-existent words (for example, *blacket*, *snord*), then participants will have to process them more deeply before deciding whether the targets are words or non-words.

masked priming A priming task in which the prime word is presented subliminally, that is, too quickly to be consciously recognized.

interstimulus interval (ISI) The amount of time between the offset of the prime and the onset of the target.

contain other parts or properties. This doesn't seem like an unreasonable assumption for recognizing written words—after all, the letter *p* looks pretty much the same, regardless of which other letters it's next to (at least if it's typed or printed). So it doesn't seem odd to think of it as a basic unit. Even so, by treating letters as integral units, the model predicts—without justification—that similarities among letters should be irrelevant. For example, it predicts that the word pair *bin*/*din* should be no more confusable than *bin*/*sin*, which is probably wrong.

But things get especially complicated in dealing with the perception of *spoken* language. Individual sounds vary a great deal, depending not only on

where in the word they appear, but also on the particular speaker's age, gender, regional dialect, and individual characteristics of the vocal tract. What's more, sounds often smear together when pronounced, making it tricky to carve up the speech stream into separate units. We'll take up these issues in Section 7.4.

Another built-in assumption that we've made so far is that word representations themselves are discrete units. Our model captures these as individual nodes that become activated or inhibited as entire units. I've been extremely vague about what's *in* a word representation. Let's assume that a word node is really just a container for information about a word's meaning and sound, as well as information specifying how the word combines syntactically with other words. But by representing words as nodes, our model implies that all of this information is available simultaneously and all to the same degree when the word node is activated. In fact, there's an ongoing debate about whether words

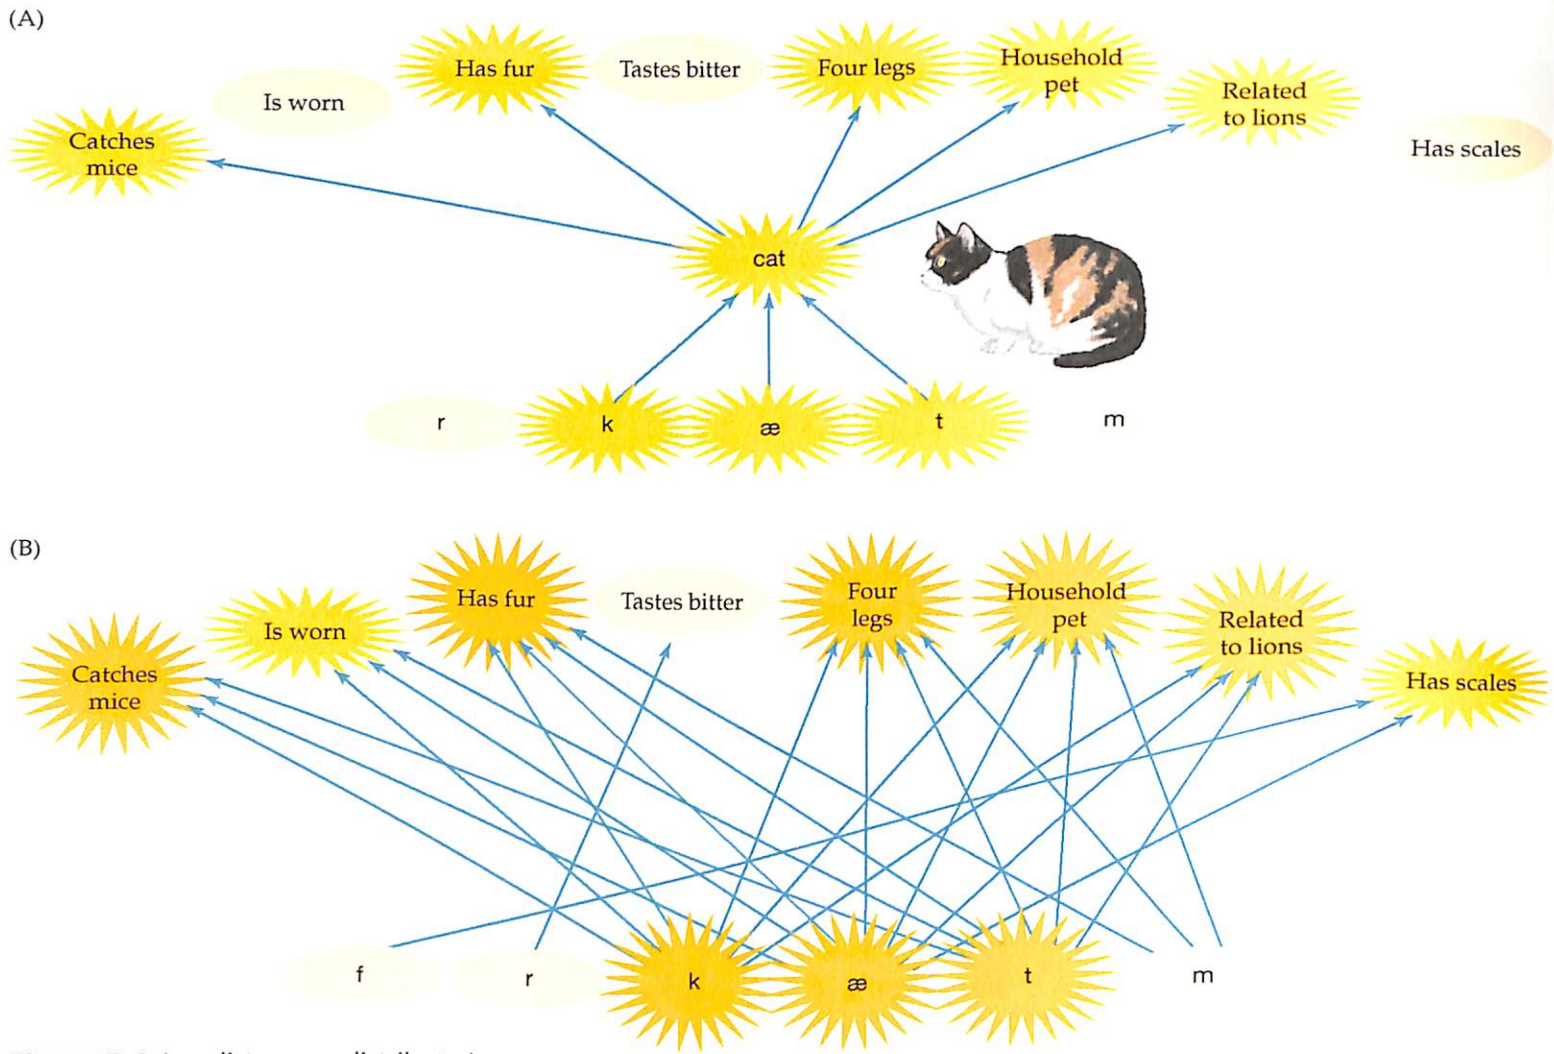

Figure 7.4 Localist versus distributed representations. (A) Localist word representations are shown by this example of a model in which the semantic features are connected to a single word unit (only excitatory connections are shown here). The activation of phonemic units results in the activation of the word unit and, in turn, the activation of that word's semantic features. (B) Distributed representations are shown in this example of a model in which phonemic units are directly linked to various semantic features without any intervening word units.

should be represented as discrete units (or containers) at all. An alternative view is that words should be captured as bundles of features instead. In dealing with the meanings of words, for example, the difference between *localist* word representations (with discrete nodes) and *distributed* representations (just bundles of features) is shown in Figure 7.4.

In the localist representation (Figure 7.4A), sound units connect to word nodes, which in turn connect to semantic features that become turned on when the word node is activated. In the distributed representation (Figure 7.4B), bundles of sound units connect directly to bundles of semantic features, without any intervening word nodes. The advantages and disadvantages of each approach, as far as generating correct predictions goes, are quite subtle, and I won't elaborate on them. But just to give you a taste for where they might differ, I'll point out that one argument that's been made for distributed representations is that they more accurately reflect some of the losses that might happen with brain damage. For example, when people suffer from aphasia due to a stroke, access to word representations is often impaired. But the damage doesn't seem to selectively wipe out swaths of a person's vocabulary, as might be expected if words were represented as integral units. Rather, the damage seems to lead to an across-the-board dysfunction in access, as if some *parts* of all word representations were destroyed, rather than some subset of word units in a person's vocabulary. (See Box 7.1 for additional discussion on how to think about word meanings.)

BOX 7.1

Words: All in the mind, or in the body too?

So far, we've been talking about meanings of words as clusters of semantic features or properties that are linked to word representations and become "turned on" when a word is heard or read. For example, the word *dog* might be represented with properties such as [has fur], [barks], [has four legs], [is a mammal], and so on. The assumption is that we've learned and mentally stored some properties about the things in the world that are called *dog*. It's often thought that these properties are abstract—that is, the property [has fur] comes to be stored in the mind in the same way regardless of how we came to know this fact about dogs, whether through visual observation of the fur, by feeling it, or by being told that dogs have fur. In other words, we don't think about meanings in terms of pictures or tactile experiences; we think about them in *thoughts*. The perceptual experiences are just the delivery device for getting to the meanings; they're not part of the meaning representation itself. To take a slightly more provocative example: the idea is that even though the color blue is closer to green on the color spectrum than it is to red, the *meaning* of the word *blue* in our minds would be no closer to the meaning of *green* than to that of *red*, since word meanings are just abstractions.

But many researchers argue that links between words and our perceptual experiences are preserved, and some even go so far as to say that meaning isn't really about pulling abstract properties from those experiences so much as it is about encoding the perceptual memories and linking them to words and to each other.

A number of studies suggest that at the very least, there is cross talk between word representations and bodily information. For example, people are usually faster to respond to a word like *pen* or *knife* when their hands are positioned as they would be if they were actually using these objects (Klatzky et al., 1989). They are faster to respond to sentences like *He closed the drawer* if the response requires them to make a movement that's consistent with simulating the action—for example, if subjects have to move the hand away from the body to push a button rather than starting at a faraway resting point and moving the hand toward the body to push a button (Glenberg & Kaschak, 2002). Semantic priming shows that reading a word like *typewriter* speeds up responses to a word like *piano*, where the main similarity between the words is in how the corresponding objects are physically manipulated, rather than some more abstract property that they share (Myung et al., 2006). And people seem to take less time to read words when they can easily imagine ways in which they might physically interact with their corresponding objects, such as *cat* versus *sun* (Phillips et al., 2012).

But it's one thing to show that words set off resonances with corresponding perceptual memories, and quite another to show that knowing and retrieving the meanings of words somehow *depends* on accessing those perceptual memories, which is what more fervent proponents of embodied meanings want to claim. It would be useful to see what happens to word processing when there is damage to sensory or motor systems in the brain. One study by Véronique Boulenger and colleagues (2008) looked at people with Parkinson's disease, a condition that dampens the brain activity responsible for planning physical movement. Would such subjects show a selective impairment in accessing the meanings of words involving action? They did; compared with other subjects, those with Parkinson's showed reduced priming for action words (suggesting that these words were being weakly activated) while showing normal priming for words that don't evoke actions. But when these subjects were treated with medication that improves the brain functioning in motor areas, they showed normal priming for action words as well. This suggests that perhaps word meanings are not entirely abstract, and that some are woven tight with bodily action. Other classes of words seem to be more tightly linked with specific sensory domains. For instance, damage to the visual region in the brain often leads to difficulty processing words for things that we normally experience visually, such as *birds* (Warrington & Shallice, 1984).

Results like these are hard to explain with a story that says word meanings are made out of pure thought rather than out of pictures or patterns of movement. They're driving a rethinking about how to talk about the meanings of words. Of course, building meanings out of body memories poses some challenges of its own, the greatest of these being how to account for more abstract words like *freedom* and *hypothesis*, let alone *if*, *but*, or *not*.

WEB ACTIVITY 7.4

Bottom-up versus top-down links In this activity, you'll work through the implications of drawing links going only from the bottom up, from the sound level to the word level, in contrast to a model that allows a bidirectional flow of information between levels.

Finally, you might have noticed that in Figure 7.3, the links between the sound units and word units aren't bidirectional; they go in one direction only, from the lower to the higher level. It certainly makes sense to have activation flow from the sound level to the word level—after all, we recognize words mostly on the basis of their sounds. But we might question whether this is the *only* direction in which information can flow. In theory, our model could have been developed otherwise, with facilitation and inhibition going from the top down as well.

I'll take this issue up in later sections, after you've had a chance to think about the empirical consequences of each move, with the help of prompts in Web Activity 7.4.

7.2 Ambiguity

A multiplicity of meanings

In the preceding section, you learned that retrieving a word from memory is not a simple matter of pulling out a single word from its designated slot in the mind's vending machine. Instead, multiple word representations are simultaneously activated, resulting in competition among lexical representations, sometimes to a degree that causes a discernible delay in the time it takes to retrieve a word. Competition is most intense when words overlap a great deal in terms of their sounds or orthography (so that *bean*, for example, might interfere with the retrieval of *beat*). Which might lead us to ask: What about when the sound overlap between words is not partial, but *complete*? As it happens, English is riddled with **homophones**, words that mean completely different things and may be spelled differently, but that sound exactly the same. Here's just a small sample:

bred, bread	made, maid	side, sighed	flea, flee
none, nun	blew, blue	missed, mist	main, mane
bridal, bridle	waste, waist	know, no	in, inn
sun, son	stare, stair	seen, scene	fair, fare
team, teem	pea, pee	hour, our	retch, wretch

The problem of ambiguity gets worse. English is also rife with **homographs**, words that share the same spelling but have different meanings (and may or may not sound the same). Consider, for example:

The performer took a deep **bow**.
It's difficult to hunt with a **bow** and arrow.

Jerry is headed **down** the wrong road.
I've really been glad to have my **down** parka this winter.

Silvia is **content** with her lot in life.
The **content** of this course is difficult.

To round the list off, consider the many words that are **polysemous**, conveying a constellation of related but different meanings. Consider some the many possible uses of the word *run*:

She's got a **run** in her stockings.

There was a **run** on the banks this week.

homophones Two or more words that have separate, non-overlapping meanings but sound exactly the same (even though they may be spelled differently).

homographs Words that are spelled exactly the same but have separate, non-overlapping meanings (and may or may not sound the same).

polysemous words Words that can convey a constellation of related, but different meanings, such as the various related meanings of *paper*, which can, among other meanings, refer to a specific material, or a news outlet.

Sam went out for an early morning **run**.

I'd like to **run** my fingers through your hair.

Let's **run** through the various options.

He's had a **run** of bad luck.

Can you **run** this over to the post office?

You can try generating a similar list for words like *paper* or *dish*—you may find yourself startled at how many different meanings or uses you can come up with.

Ambiguity is so rampant in language (and not just in English) that you might begin to wonder whether it's a serious design flaw common to many languages. It's hard to imagine that the presence of ambiguity does anything useful to promote effective communication between people (but see Box 7.2).

BOX 7.2

Why do languages tolerate ambiguity?

If the goal of language is for a speaker to plant his intended meaning firmly and decisively into the mind of the hearer, ambiguity appears to be a serious flaw. By definition, an ambiguous word or phrase is compatible with multiple interpretations, not just the meaning the speaker intended. You'd think that languages would strive to avoid ambiguity. Yet all known languages seem to be rife with it, despite the fact that lexical ambiguity could very easily be avoided. In the words of blogger Geoff Pullum (2012):

> *Let me make a numerical point to begin with. The number of [possible letter sequences] with length not more than 10 over the roman letters a to z plus the apostrophe is 27^{10} = 205,891,132,094,649—about 200 trillion. The total number of words in the workaday word list is about 25,000. What I'm saying is that English could easily have a distinct letter sequence for every different meaning, using letter sequences much shorter than the present ones. It doesn't because the language in general shows no signs of being the slightest bit interested in that. English uses the same two-word phrase for denigrating, ceasing to hold, making notes, and euthanasia. [The phrase is* ***put down****.] It wantonly employs a single three-letter word for meanings relating to understanding, judging, experiencing, finding out, dating, visiting, ensuring, escorting, and saying farewell. [The word is* ***see****; see if you can create a sentence for each of these uses.] Nobody who thinks about English for a few seconds could possibly believe it shuns ambiguity. It doesn't give a monkey's fart about avoiding ambiguity.*

Steve Piantadosi and his colleagues (2012) have gone even further and argued that not only do languages not "care" about avoiding ambiguity, they actively seek it out because ambiguity actually makes a language more effective. The logic goes like this: Ambiguity rarely creates serious impediments to understanding—yes, processing ambiguous words comes with a small cost for the hearer (as you'll see in the rest of this chapter), and yes, occasionally, communication may rupture as a result. But the vast majority of the time, hearers are quite competent at relying on context to navigate through the various meanings offered up by a single string of sounds or letters. The benefits of ambiguity come from considering the costs of *producing* language. Speakers can minimize their effort by re-using bits of language that are common, short, and easy to pronounce, rather than resorting to longer words with unusual combinations of sounds. The idea is that languages tend to strike a balance between comprehensibility and ease of production. If ambiguity is managed fairly easily by the hearer, the speaker may as well take advantage of it to reduce his own cognitive workload.

Throughout this chapter, you'll see how the word recognition mechanism is set up to avoid ambiguity sinkholes, suggesting that Piantadosi and his colleagues are right that ambiguity doesn't do much damage to understanding. In support of their second point—that ambiguity makes the task of speaking easier—Piantadosi and his colleagues presented evidence from several languages showing that words that are easier to produce are exactly the ones that are most likely to be re-used for new meanings. That is, ambiguous words in those languages were generally shorter, more common, and composed of fewer unusual combinations of sounds than unambiguous words.

The English writer Virginia Woolf had an interesting perspective on ambiguity and the usefulness of language. In her 1937 essay on writing titled "Craftsmanship," Woolf argued that if we think of a useful statement as one that can mean only one thing—that is, a statement that unambiguously communicates a very specific idea—then it should be apparent "how very little natural gift words have for being useful. ... They have so often proved that they hate being useful, that it is in their natures not to express one simple statement but a thousand possibilities." To make her point, Woolf suggested the reader imagine what's going on inside the mind upon hearing a simple and seemingly utilitarian phrase:

> Take the simple sentence "Passing Russell Square." That proved useless because besides the surface meaning, it contained so many sunken meanings. The word "passing" suggested the transiency of things, the passing of time, and the changes of human life. Then the word "Russell" suggested the rustling of leaves and the skirt on a polished floor also the ducal house of Bedford and half the history of England. Finally, the word "Square" brings in the sight, the shape of an actual square combined with some visual suggestion of the stark angularity of stucco. Thus one sentence of the simplest kind rouses the imagination, the memory, the eye and the ear—all combine in reading it.

For Woolf, the point of language is just that—to rouse the imagination, rather than communicate a specific idea. Poets (and maybe even advertising copywriters) would likely agree with that, and I myself would admit that it's certainly the aim of great *writing* to rouse the imagination. But I strongly doubt that whoever announced "Passing Russell Square" on the train did so with the intent of evoking the rustling of long skirts on polished floors. And I also suspect that your average passenger understood the announcer to be communicating a specific idea (even if thoughts of rustling skirts *did* enter his mind).

But Woolf's discussion of language, fanciful though it may be, is in some ways a plausible psycholinguistic hypothesis about word recognition. There is good reason to think that during the course of recognizing a single word, a plethora of meanings presents itself.

Look back at the model we built in Figure 7.3, and notice the connections between the units of sound and word representation units. So far, we've been assuming that word units are activated to the extent that the sound units they are connected to become activated; perfect matches will be activated the most, but highly similar words will also light up. But as Virginia Woolf notes, often a set of sounds will match up *exactly* with a number of different word meanings. For example, the sounds in the name *Russell* match up with the name of whoever Russell Square happens to be honoring, but also with many different Russells, as well as the homophonous word *rustle*. Presumably, the mental lexicon includes connections from this set of sounds to *all* of these word representations, so activation of the sound units should spread to all of these word representations and their respective semantic features. If that's the case—and if, as suggested by Virginia Woolf, words commonly have a multiplicity of meanings—then the real puzzle is: How is it that words manage to make themselves useful after all? In other words, how is it that we ultimately arrive at a single interpretation, despite the numerous possible meanings?

The easy answer to this question is that we undoubtedly use context to disambiguate meanings. But since we've been building detailed models, let's be a bit more precise: *How* do we use context? We might build up our model using two different approaches.

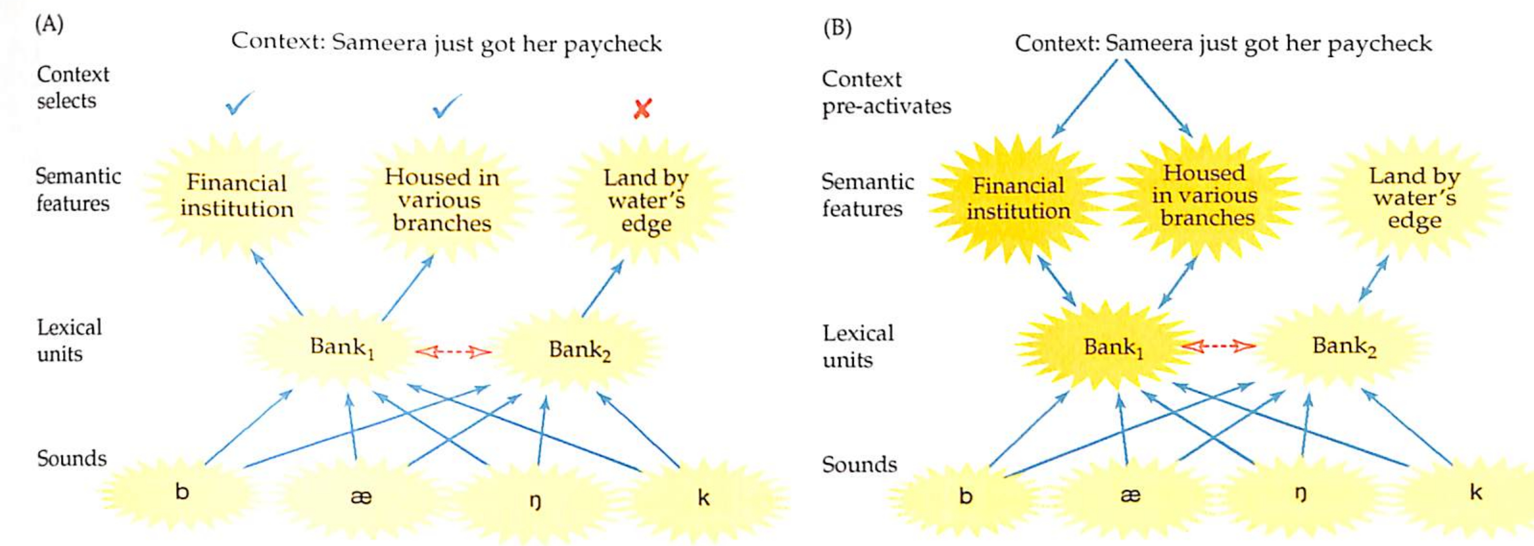

Figure 7.5 Two ways in which context can help word recognition. (A) Activation flows from the bottom up, from phonemic units to words and in turn to semantic features. Both meanings of the word *bank* are equally active until contextual information is recruited to select the most appropriate meaning. (B) Context can generate more expectation for some meanings than others by "pre-activating" some semantic features. Hence, by the time the word *bank* rolls off the tongue of the speaker, one of the meanings of that word is already more active than the other. Inhibitory connections appear in red.

The first is illustrated in **Figure 7.5A**, using two meanings of the word *bank*. Here, sound units are activated as these sounds are heard. Activation flows from the sound level to the lexical level, activating both word representations and, in turn, their associated semantic features. The flow of activation is in one direction only, from the lower level of sound to the higher level of meaning. (Lateral inhibitory links have also been drawn in between word representations, but since each word representation receives equal activation, these links would have no impact on the relative activation of the two possible meanings.) Once words and their meanings are activated, a separate decision mechanism is triggered to select the most contextually appropriate word, based on a good match between the semantic features of that word and the semantic expectations that have been set up by the context. To relate this to Virginia Woolf's observations, this would mean that multiple meanings do in fact routinely flare up in the mind, even if we ultimately have a way of picking out the most "useful" one.

A second approach is shown in **Figure 7.5B**. The crucial difference here is that activation can move from the top semantic level down to lower levels. Context activates certain semantic features, which in turn activate associated word representations. This means that even before there's any sound input, one word representation may be more active than the other. Once the word itself is uttered, activation moves from the sound units to each of the competing meanings, but since the contextually favored word is already more strongly activated, it inhibits the less-favored reading. As a result, the activation level of the competing meaning may remain very weak, perhaps even negligible. Our response to Virginia Woolf might be: even though language is rife with ambiguity, the mind is very efficiently set up to promote the most "useful" meanings of words at the expense of the less useful ones. In that case, what do we make of Woolf's poetic meditations about rustling skirts? Well, perhaps such alternative meanings come to mind when you think about language in a more deliberate way, with a purposeful focus on the connections between words. But they're unlikely to spontaneously arise in the normal course of language comprehension.

Evidence for the simultaneous activation of word meanings

In 1979, David Swinney published a seminal study that tested whether competing meanings of ambiguous words become simultaneously activated even when plenty of contextual information makes it easy to home in on the most useful meaning. For example, the word *bug* could refer to either a small crawl-

crossmodal priming task An experimental task involving both spoken and written modalities; participants typically hear prime words, which are often embedded within full sentences, and they must respond to test words displayed orthographically on a computer screen.

ing creature or a surveillance device for eavesdropping, but its meaning should be clear in the following context:

> Rumor had it that, for years, the government building had been plagued with problems. The man was not surprised when he found several spiders, roaches, and other bugs in the corner of his room.

Swinney devised an elegant experiment to test the following predictions: Clearly, the relevant meaning of the word *bugs* must become activated in order for someone to understand the passage. Hence, its activation should spread to other words that are related in meaning—for example, the word *ant*. Now, if the irrelevant meaning *also* becomes activated, then its activation should spread as well to related words, such as *spy*. But if the activation of the irrelevant meaning remains low, then the activation flowing to *spy* should be no greater than to a word completely unrelated to either meaning (such as *sew*). And how to measure the spread of activation? Through the familiar semantic priming task.

Swinney used a **crossmodal** variant of the priming task in which subjects listened to passages like the one above and then responded to test words presented visually on a screen (see Figure 7.6). Subjects responded by pressing one of two buttons to indicate whether they thought they'd seen a word or a non-word (naturally, plenty of non-word filler items were included). Subjects saw the test word *ant*, *spy*, or *sew*. Swinney also varied whether subjects heard passages with helpful context or with context that didn't help to disambiguate between the meanings of *bugs*. To complete the comparison, he also varied whether subjects heard the ambiguous word *bugs* in the passage, or an unambiguous counterpart, *insects*.

Figure 7.6 The crossmodal priming task. (A) Sample experimental materials for the task as used by Swinney (1979). (B) A subject from one of the four experimental conditions listens to sentences while sitting in front of a computer screen. At some predetermined point in the sentence, a string of letters appears on the screen, and the subject must press a button to indicate whether the letter string is a real word or a non-word. (Adapted from Swinney, 1979.)

(A)

BIASING CONTEXT

Condition 1: Ambiguous prime
"Rumor had it that, for years, the government building had been plagued with problems. The man was not surprised when he found several spiders, roaches, and other bugs in the corner of his room."

Condition 2: Unambiguous prime
"Rumor had it that, for years, the government building had been plagued with problems. The man was not surprised when he found several spiders, roaches, and other insects in the corner of his room."

NEUTRAL CONTEXT

Condition 3: Ambiguous prime
"Rumor had it that, for years, the government building had been plagued with problems. The man was not surprised when he found several bugs in the corner of his room."

Condition 4: Unambiguous prime
"Rumor had it that, for years, the government building had been plagued with problems. The man was not surprised when he found several insects in the corner of his room."

(B)

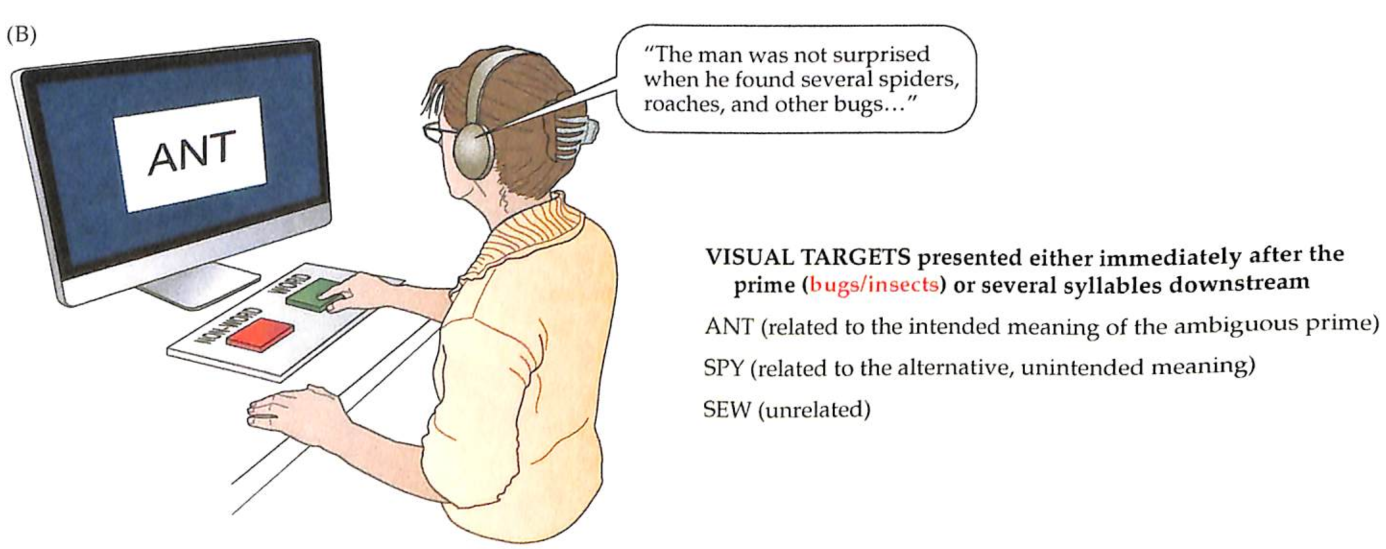

When subjects had to respond to the target word right after hearing the ambiguous word *bugs*, Swinney saw clear evidence of the activation of *both* meanings of that word. That is, response times were faster to both *ant* and *spy* than they were to the control word *sew*. In contrast, the pattern was quite different when the unambiguous word *insects* was substituted for *bugs*: here, as would be expected, response times for *ant* were faster than for *sew*, but response times for *spy* were not. This shows that it was clearly the presence of the ambiguous word *bugs* that caused the activation of the word *spy*. So Virginia Woolf is at least partly right: even the useless "sunken meanings" become active in the mind along with the obviously intended "surface meanings."

But Swinney's experiment also showed that people quickly converge on the intended meaning. When the same experiment was repeated but with the test word appearing on the screen three syllables after subjects heard the word *bugs* (or *insects*), the results were quite different. In that case, responses to the word *ant*, but not to the word *spy*, were sped up relative to *sew*. It appears that the irrelevant meaning had been activated in parallel with the more pertinent one, but then was quickly suppressed.

These results support a model in which all lexical candidates that match the sound input become active, at least for a time. But later research suggests that the picture is actually a bit more complex, and that in some cases, contextually inappropriate meanings never reach discernible levels of activation. The full story, it turns out, needs to take into account the relative frequencies of the competing meanings.

To see how frequency of meanings might play a role, play along with the following exercise. Quick, now: for each of the words below, say aloud the first synonym or brief definition that comes to mind:

port	straw	chest	pitcher
yarn	cabinet	bark	mint

All of these words are ambiguous. But for some of them, one particular meaning may have sprung to mind immediately, while for others you may have been aware of a mental tug-of-war between two meanings. And if you compared your answers with your those of classmates, some words would have converged on one meaning, while others may have been split between them. This is because for some words (for example, *bark, pitcher, straw, chest*) both meanings occur with roughly equal frequency in English, while for others (for example, *port, cabinet, yarn, mint*) one meaning is dominant and the other is subordinate. A well-known study by Susan Duffy, Robin Morris, and Keith Rayner (1988) explored how this factor interacts with contextual expectations. Rather than using semantic priming as the basis for their experiment, they exploited another behavioral consequence of lexical activation: remember that when multiple words are activated at the same time and compete with each other, as in the studies of neighborhood density effects, people usually take longer to recognize the target word. As you might expect, ambiguous words generally take longer to read than unambiguous ones, providing some additional evidence that multiple meanings of these words are activated simultaneously and compete with each other. And, based on David Swinney's priming results, we might predict that such competition would show up regardless of whether the context favors one of these readings or not, since both meanings appear to be activated, at least for some time.

But Duffy and colleagues found something a bit more subtle. When the context favored the subordinate meanings of words like *mint* or *cabinet*, subjects did read these more slowly than unambiguous control words in the same sentence, indicating competition from the alternative meanings. (For example *mint* was read more slowly than *jail* in a sentence like *Although it was by far the largest building in town, the mint/jail was seldom mentioned.*) But when the words were equally biased in frequency between the two meanings, as for *pitcher* or *straw*, and when the context favored one of their meanings, people spent no more time reading these ambiguous words than unambiguous control words, suggesting that there was no discernible competition from the alternative meaning. This raises the possibility that both frequency of meanings *and* contextual expectations can affect the activation levels of word representations. When these conspire to boost the activation of the same meaning, this lexical representation becomes disproportionately activated, allowing it to very quickly inhibit any potential competitors; however, when the two sources of information conflict, with frequency boosting the activation of one and context favoring the other, the result is roughly equal activation of both meanings, leading to competition between them.

LANGUAGE AT LARGE 7.1

The persuasive power of word associations

What if hearing a word didn't set off resonances within an entire connected network of information? What if word recognition really did just work like a vending machine, with the sounds and letters of words merely acting as pointers or addresses to slots containing meanings or concepts?

If that were the case, marketers, advertisers, and politicians would probably be a lot less preoccupied with words than they are. For example, in a word-vending-machine world, a politician trying to persuade the public on a point of policy shouldn't really care what that policy is called. After all, different ways of saying the same thing would just be different ways to get to the same concept slot in the vending machine.

But politicians and their strategists do care about the names for policies, sometimes obsessively so. For example, Frank Luntz (2007), who has worked as a communications consultant for Republican party candidates in the U.S., has a list of suggestions, seen in the table, for the wording of various policies or initiatives to advance the interests of his clients.

Never say	Instead, say
Tax cuts	Tax relief
Drilling for oil	Energy exploration
Private health care	Free market health care
Wiretapping	Electronic intercepts
Estate tax	Death tax

If you feel insulted at the suggestion that you might feel differently about a plan to reduce taxes depending on the name used to refer to it, you might consider the results of a 2010 poll conducted by CBS and the *New York Times* to probe how Americans felt about having gay people serve in the U.S. military. When people were asked, "Do you favor or oppose gay men and lesbians serving in the military?" the results were as follows:

Strongly favor: 51%

Somewhat favor: 19%

Somewhat oppose: 7%

Strongly oppose: 12%

But when people were asked, "Do you favor or oppose homosexuals serving in the military?" they were less receptive to the idea:

Strongly favor: 34%

Somewhat favor: 25%

Somewhat oppose: 10%

Strongly oppose: 19%

This makes no sense if words are just pointers to concepts. But as you now know, an awful lot happens in the mind on the way to retrieving a word's meaning from memory. Different patterns of activation will resonate throughout the lexicon, depending on whether you hear a phrase like *drilling for oil* or *energy exploration*. (If you want to get a feeling for the differences, you might try a little experiment: give two groups of people a word association task like the one you did in Web Activity 7.1. Among the words on the list, include *energy*, *oil*, *exploration*, and *drilling*, and see what comes up.)

A growing body of evidence suggests that the activations that are set off during word recognition probably amount to more than just brief mental flickers that quickly dissipate without any consequences for behavior. Within social psychology, researchers have studied a phenomenon known as **implicit priming**, in which exposing people to certain stimuli increases the likelihood that they'll behave in ways that reflect stored associations, which are activated upon perceiving the stimuli. For example, in one classic study by John Bargh and colleagues (1996), undergraduate students formed sentences out of scrambled word lists, with some students receiving lists that contained words associated with the elderly (for example, *Florida*, *wrinkles*, *bingo*, *gray*) while other students got lists of neutral control words. After the students had finished the test, the experimenters measured and compared how quickly students from the two groups walked down the hall. Those who'd been exposed to the words associated with the elderly walked more slowly than those who'd been in the control condition.

Naturally, marketers are also highly intrigued by the possibility of meaningful links between word associations and behavior or attitudes, and such links are increasingly being tested in the lab by researchers who are interested in the psychology of consumer behavior. To give you just one example, Jonah Berger and Gráinne Fitzsimons (2008) constructed an experiment to see whether exposing people to photographs of dogs would make them more likely to give positive evaluations of sneakers carrying the brand name of Puma. In case the logic behind this study escapes you, it goes like this: Generally, the more familiar people are with an idea or concept, the more they're inclined to like it (which explains why you might have the same TV commercial inflicted upon you half a dozen times during a single program). Berger and Fitzsimons reasoned that because of the similarity between the concepts of dog and puma, pictures of the dog would activate the Puma brand name, making it feel more familiar. As a result, people should experience warmer feelings toward the Puma products, which, in fact, was what the researchers found.

Naturally, just because you see an effect in the lab doesn't mean it will carry the day out in the real world. Real-world choices made by consumers or voters are complex, and likely to be affected by a wide range of different variables; I seriously doubt that you'd be convinced to buy a product you otherwise have no interest in, simply because of the associations set off by its name. Nevertheless, experimental research does lend some credibility to the notion that the Edsel automobile, one of the greatest marketing flops of the last century, wasn't helped in any way by its name. The car was named after one of its makers, Edsel Ford, who unfortunately bore a highly unpopular, old-fashioned, Germanic-sounding name—perhaps not the right one to attach to an American car a mere decade or so after the Second World War.

implicit priming A psychological phenomenon in which exposing people to certain stimuli increases the likelihood that they'll exhibit behaviors that are associated with the stimuli. For example, exposing people to words associated with the elderly may trigger behaviors that are stereotypically associated with the elderly, such as walking slowly.

These additional results make it clear that both of the models in Figure 7.5 need to be adjusted to take into account the effects of frequency and the way it interacts with context. But *exactly* how these factors interact is still a matter of some debate. (For example, do they both exert an influence at exactly the same point in time during word recognition, or does one factor come into play earlier than the other?) As a result, researchers are still in the process of refining the models to greater and greater levels of precision and continue to test finer and finer predictions about the behaviors that should result. But there's general agreement that multiple sunken meanings *are* often aroused in the mind (even though the ambiguity doesn't generally impede the eventuality of getting to a single useful meaning), and that the extent to which this occurs depends jointly on the context and frequency of the alternative meanings.

Virginia Woolf's literary excursions raise some questions that don't seem to come up in the more scientific literature. For the most part, even when there is

WEB ACTIVITY 7.5

Product names and sunken meanings Many product names are deliberately chosen from among the inventory of existing English words (for example, the product names *Apple* and *Tide*). This creates a new lexical ambiguity, where the word can now refer either to its original meaning, or to the newly named product. In this activity, you'll explore some possible implications of this practice.

clear evidence of parallel access of competing meanings, the irrelevant meaning is quite fleeting, and quickly submerged. But do these active meanings, however fleeting, nevertheless manage to have an impact on our aesthetic or emotional experience of language? We don't really know—but recent findings and discussions about how alternative ways of saying the same thing may have different persuasive effects (see Language at Large 7.1) suggest that, perhaps, even brief flickers of activation from "useless" meanings or associated words may not be inconsequential.

7.3 Recognizing Spoken Words in Real Time

The flow of spoken words

So far, most of what we've seen about word recognition could apply equally well to the spoken or written language modality, and in fact, the experimental methods that we've seen have relied on both spoken and written stimuli, and sometimes both, to explore the underlying psychological mechanisms. But spoken language offers some particular challenges for hearers, along with some specific puzzles for researchers. We now turn to these modality-specific issues.

One obvious difference between spoken and written words is that when you read a word on a screen or on the page, you can see the whole word at once, and in normal circumstances, you can stare at it for as long as it takes to recognize it. But spoken language unfolds one sound at a time, rather than being uttered all at once, and once it's been uttered, it's gone. As aptly described in a paper by James Magnuson and colleagues (2007), if reading were like listening to spoken language, it would be like this: "Imagine reading this page through a two-letter aperture as the text scrolled past, without spaces separating words, at a variable rate you could not control." It would feel deeply weird to read a word that appeared one or two letters at a time from left to right, and this intuition is confirmed by studies that look at where people focus their gaze while reading. Rather than scanning the word left to right, their gaze lands somewhere within the word, and they can usually read the entire word from that position (or, if the word is very long, they might move their eyes rightward once to take in the rest of the word).

The "scrolling by" nature of spoken word recognition raises a very interesting question: At what point do people initiate the process of matching a string of speech sounds to a stored word representation? In the previous sections, you saw how in the process of word recognition, activation flows from sounds to word representations that contain those sounds. For example, when you hear the sequence of sounds /k/, /ae/, and /t/, the word *cat* will be activated, and also, to a lesser extent, the words *cot* and *can*, among others. But when does the activation of possible word representations begin, given that there's a time lag between the first and last sounds of a word? Does the activation of word candidates start even before the end of the word is encountered in the speech stream, or is it delayed until all the sounds of the word have been uttered? And if people do wait until the entire word has been uttered before activating lexical candidates, how do they identify where the end of the word is anyway, given that usually no silences occur between words in running speech?

In principle, it should be perfectly possible to first locate likely word boundaries, and then activate all of the sounds that bundle together in one word so that they in turn can activate matching lexical candidates. Remember, after all, that we saw in Chapter 4 that even tiny babies were able to figure out where to break the speech apart into words on the basis of statistical information, probably months before they'd acquired much in the way of a working lexicon. So one could imagine that adult word recognition might rely on the same kind of statistically based word segmentation, which would serve as the very first step in spoken word recognition—in a way, we'd be mentally inserting "spaces" between the words before any lexical activation occurred. An analogy with text might be that you'd run a program, based on statistical probabilities, to put spaces between the spoken words before any actual "reading" of the words themselves began.

How can we test to see whether lexical activation begins only after both edges of the word are identified, or whether it's initiated before this point? We can readily recruit some of the methods we've already talked about, and simply tweak them a bit.

For example, remember that in his famous study, David Swinney used semantic priming as a tool to probe for the level of activation of competing word representations: evidence of priming (that is, speeded responses) for words like *spy* and *ant* meant that both meaning representations for *bugs* were activated. We can use a similar logic now, but instead of giving subjects entire words, we can present *partial* words and see whether there's any evidence of priming for words that are related to potential matches to partial words. For example, imagine recording a word like *conform*, and cutting off the sound file right in the middle of the sound /f/. Statistically, the sound sequence /nf/ is extremely unlikely to correspond to the end of a word, so hearers should be able to guess that the end of the word hasn't occurred yet, based solely on information about the sound patterns of English words. So, if lexical activation is delayed until the ends of words are identified, we wouldn't expect to see priming for any words that are related to *conform*—for example *copy* or *imitate*. On the other hand, if lexical activation *is* initiated, we'd expect to see priming not only for words related to *conform*, but also for words related to other possible continuations of this snippet, that is, words semantically related to *conflate*, *confabulate*, *confuse*, *confine*, *confide*, *conflicted*, and so on. Such words, with their overlapping onsets, are known as **cohort competitors**.

This latter scenario is exactly what William Marslen-Wilson (1987) predicted. In his **cohort model** of word recognition, he suggested that lexical activation begins right after the beginning of a word, with multiple cohort competitors becoming active. As more and more sound input comes in over time as the word unfolds, the set of possible matching candidates dwindles until the **uniqueness point**, at which there remains only one possible match with the sound input. Table 7.1 illustrates how the set of cohort candidates becomes smaller and smaller with each incoming snippet of speech.

cohort competitors Words with overlapping onsets (e.g., *candle*, *candy*, *candid*, etc.).

cohort model A model of word recognition in which multiple cohort competitors become active immediately after the beginning of word is detected, and are gradually winnowed down to a single candidate as additional acoustic information is taken in.

uniqueness point The point at which there is enough information in the incoming speech stream to allow the hearer to differentiate a single word candidate from its cohort competitors.

TABLE 7.1 Winnowing down cohort candidates as a word unfolds in time

Initial sounds heard	Cohort candidates
/kæ /	*cat, cap, cast, can, cash, cad, camp, cab, cattle, capture, candidate, catholic, candelabra, captain, canteen, castrate, Canada, cancel, castle, canister, captive, candle, cantaloupe, castoff, candy, cannibal, cashew, cantankerous, California, castaway*, many others
/kæn/	*can, candidate, candelabra, canteen, Canada, cancel, canister, candle, cantaloupe, candy, cannibal, cantankerous*, others
/kænɪ/	*canister, cannibal*
/kænɪs/	*canister*