Language

10

What Is Language?
- The Creativity of Human Language
- The Universality of Language
- Studying Language in Cognitive Psychology

Perceiving and Understanding Words
- Components of Words
- Perceiving Words
- Understanding Words
- Demonstration: Lexical Decision Task
- Method: Lexical Priming
- Summary: Words Alone and in Sentences

Understanding Sentences
- Semantics and Syntax
- Method: Event-Related Potential
- Parsing a Sentence

Test Yourself 10.1

Understanding Text and Stories
- How Inference Creates Coherence
- Situation Models

Producing Language: Speech Errors
- Method: Identifying Speech Errors
- Speech Errors and Language Mechanisms

Producing Language: Conversations
- Semantic Coordination
- Syntactic Coordination
- Method: Syntactic Priming

Something to Consider: Culture, Language, and Cognition
- Method: Categorical Perception

Test Yourself 10.2

Chapter Summary

Think About It

If You Want to Know More

Key Terms

CogLab: Word Superiority; Lexical Decision; Categorical Perception—Identification; Categorical Perception—Discrimination

Some Questions We Will Consider

- How do we understand individual words, and how are words combined to create sentences? (363)
- How can we understand sentences that have more than one meaning? (370)
- What do speech errors (slips of the tongue) tell us about language? (381)
- Is it true that the language that people use in a particular culture can affect the way they think? (387)

There are ways to communicate that don't involve language, but language is the most powerful tool we have for transmitting ideas, feelings, and knowledge from one person to another. What exactly is language, and what is it about language that makes it so useful?

What Is Language?

We can define **language** as *a system of communication using sounds or symbols that enables us to express our feelings, thoughts, ideas, and experiences.* Although one of the main features of language is communication, it is important to differentiate human language from the communication of nonhuman animals. Cats "meow" when their food dish is empty. Monkeys have a repertoire of "calls" that stand for things such as "danger" or "greeting," and bees signal through a "waggle dance" that they perform at the hive to indicate the location of flowers. Although there is some evidence that monkeys may be able to use language in a way similar to humans (see "If You Want to Know More: Animal Language"), most animal communication lacks the properties that make human language unique.

The Creativity of Human Language

Human language goes far beyond a series of fixed signals that transmit a single message like "feed me," "danger," or "go that way for flowers." Language provides a way of arranging a sequence of signals—sounds for spoken language, letters and written words for written language, and physical signals for sign language—that provide a wide variety of ways to transmit, from one person to another, things ranging from the simple and commonplace ("My car is over there") to things that have perhaps never been previously written or uttered in the entire history of the world ("I'm thinking of getting a new Mustang because I'm quitting my job in February and taking a trip across the country to celebrate Groundhog Day with my cousin Zelda").

Language makes it possible to create new and unique sentences because it has a structure that is (1) hierarchical and (2) governed by rules. Language is hierarchical because it consists of a series of components that can be combined to form larger units. For example, words can be combined to create phrases, which, in turn, can create sentences, which themselves can become components of a story. Language is governed by rules that specify permissible ways for these components to be arranged ("What is my cat saying?" is permissible in English; "Cat my saying is what?" is not). These two properties—a hierarchical structure and rules—endow humans with the ability to go far beyond the fixed calls and signs of animals to communicate whatever they want to express.

The Universality of Language

Although people do "talk" to themselves, as when Hamlet wondered "To be or not to be" or when you daydream in class, the predominant staging ground for language is one person conversing with another. Consider the following:

- People's need to communicate is so powerful that when deaf children find themselves in an environment where there are no people who speak or use sign language, they invent a sign language themselves (Goldwin-Meadow, 1982).
- Everyone with normal capacities develops a language and learns to follow its complex rules, even though they are usually not aware of these rules. Although many people find the study of grammar to be very difficult, they have no trouble using language.
- Language is universal across cultures. There are over 5,000 different languages, and there isn't a single culture that is without language. When European explorers first set forth in New Guinea, the people they discovered, who had been isolated from the rest of the world for eons, had developed over 750 different languages, many of them quite different from one another.
- Language development is similar across cultures. No matter what the culture, children generally begin babbling at about 7 months, a few meaningful words appear by the first birthday, and the first multiword utterances occur at about age 2 (Levelt, 2001).
- Even though a large number of languages are very different from one another, we can describe them as being "unique but the same." They are unique because they use different words and sounds, and they may use different rules of combining these words (although many languages use similar rules). They are the same because all languages have words that serve the function of nouns and verbs, and all languages include a system to make things negative, to ask questions, and to refer to the past and present.

Studying Language in Cognitive Psychology

Wilhelm Wundt, founder of the first laboratory of scientific psychology, wrote about the nature of the sentence in 1900, but as with other areas of cognitive psychology, modern research on language had to await the "cognitive revolution" that began in the 1950s. Two events that occurred during that time stand out. The first was the 1957 publication of a book by B. F. Skinner, the modern champion of behaviorism. In this book, *Verbal Behavior*, Skinner proposed that language is learned through reinforcement. According to this idea, just as children learn appropriate behavior by being rewarded for "good" behavior and punished for "bad" behavior, children learn language by being rewarded for using correct language and punished (or not rewarded) for using incorrect language.

In the same year, the linguist Noam Chomsky published a book titled *Syntactic Structures.* This book, and Chomsky's work that followed, proposed that human language was coded in the genes. According to this idea, just as humans are genetically programmed to walk, they are programmed to acquire and use language. Chomsky concluded that despite the wide variations that exist across languages, the underlying basis of all language is similar. Most important for our purposes, Chomsky saw studying language as a way to study the properties of the mind and therefore disagreed with the behaviorist idea that the mind is not a valid topic of study for psychology.

Chomsky's disagreement with behaviorism led him to publish a scathing review of Skinner's *Verbal Behavior* in 1959. In his review, he presented arguments that effectively destroyed the behaviorist idea that language can be explained in terms of reinforcements and without reference to the mind. One of Chomsky's most persuasive arguments was that as children learn language, they produce sentences that they have never heard and that have therefore never been reinforced. (A classic example of a sentence that has been created by many children, and which is unlikely to have been taught by parents, is "I hate you, Mommy.") Chomsky's criticism of behaviorism was one of the most important events of the cognitive revolution and led to the development of **psycholinguistics**, the field concerned with the psychological study of language.

The goal of psycholinguistics is to discover the psychological processes by which humans acquire and process language (Clark & Van der Wege, 2002; Gleason & Ratner, 1998). The three major concerns of psycholinguistics are as follows:

1. *Comprehension.* How do people understand spoken and written language? This includes how people process language sounds; how they understand words, sentences, and stories, as expressed in writing, speech, or sign language; and how people have conversations with one another.
2. *Speech production.* How do people produce language? This includes the physical processes of speech production and the mental processes that occur as a person creates speech.
3. *Acquisition.* How do people learn language? This includes not only how children learn language, but also how people learn additional languages, either as children or later in life.

Because of the vast scope of psycholinguistics, we are going to restrict our attention to the first two of these concerns by describing research on how we understand language and how we produce it. (See "If You Want to Know More: Language Acquisition" for suggestions for readings about language acquisition.) We begin by considering each of the components of language, beginning with small components such as *sounds* and *words* (Figure 10.1), then combinations of words that form *sentences,* and finally "texts"—*stories* that are created by combining a number of sentences. At the end of the chapter, we de-

■ **Figure 10.1** Flow diagram for this chapter.

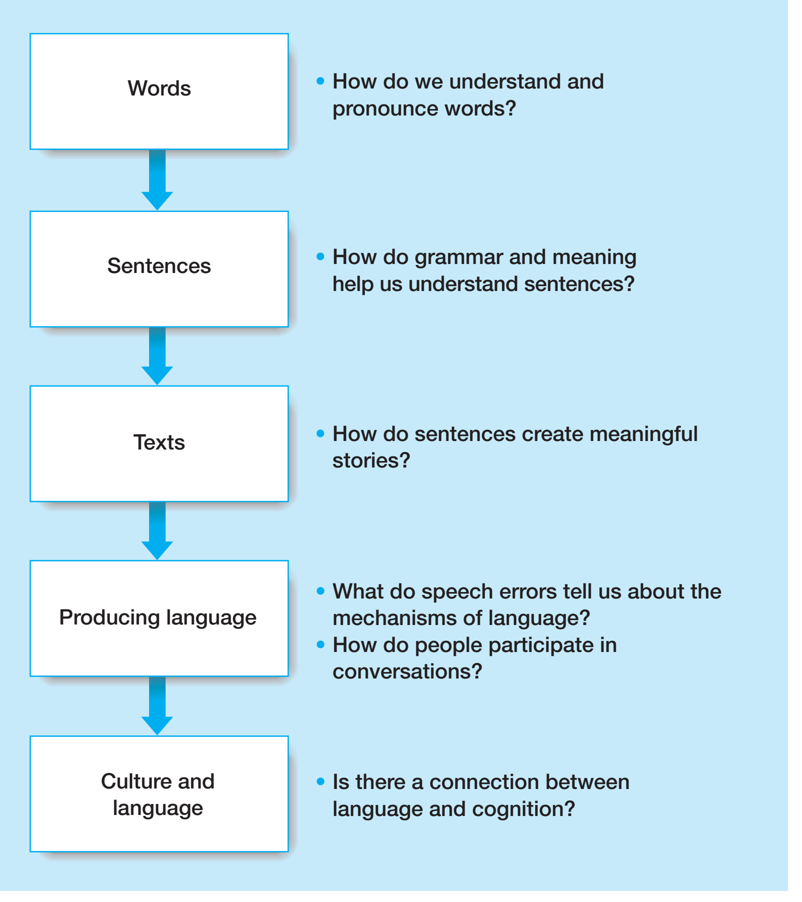

scribe some of the factors involved in producing language, considering both the errors people make while speaking and how people participate in and understand conversations. Finally, we look at cross-cultural research on language that considers the role of language in thinking.

Perceiving and Understanding Words

One of the most amazing things about words is how many we know and how rapidly we acquire them. Infants produce their first words during their second year (sometimes a little earlier, sometimes later), and after a slow start, begin adding words rapidly until, by the time they have become adults, they can understand over 50,000 different words (Altmann, 2001; Dell, 1995). All of the words a person understands are called person's **lexicon**.

Components of Words

The two smallest units of language are phonemes, which refer to sounds, and morphemes, which refer to meanings.

Phonemes

Each word you are reading is made up of letters. If you were to read these words out loud, you would produce sounds called phonemes, where a **phoneme** is the shortest segment of speech that, if changed, changes the meaning of a word. Thus, the word *bit* contains the phonemes /b/, /i/, and /t/ (phonemes are indicated by phonetic symbols that are set off with slashes), because we can change *bit* into *pit* by replacing /b/ with /p/, to *bat* by replacing /i/ with /ae/, or to *bid* by replacing /t/ with /d/.

Note that because phonemes refer to sounds, they are not the same as letters, which can have a number of different sounds (consider the "e" sound in "we" and "wet"), and which can be silent in certain situations (the "e" in "some"). Because different languages use different sounds, the number of phonemes varies in different languages. There are only 11 phonemes in Hawaiian, about 47 in English, and as many as 60 in some African dialects.

Morphemes

Morphemes are the smallest units of language that have a definable meaning or a grammatical function. "Truck" consists of a single morpheme, and even though "table" has two syllables, it also consists of a single morpheme, because the syllables alone have no meaning. In contrast "bedroom" has two syllables and two morphemes, "bed" and "room." Endings such as "s" and "ed," which contribute to the meaning of a word, are morphemes. Thus even though "trucks" has just one syllable, it consists of two morphemes, "truck" and "s."

Perceiving Words

CogLab
Word Superiority

How we perceive the letters that make up written words and the sounds that create spoken words is a huge topic. We know from our discussion of the word superiority effect in Chapter 3 that a word's meaning helps a person perceive the letters that make up the word (see page 64). The meanings associated with words create a context that makes perception of the word's components easier. The meaning of words also helps us hear a word's phonemes, even when these phonemes are obscured by another sound.

Phonemic Restoration Effect

Richard Warren (1970) demonstrated the effect of meaning on the perception of phonemes. Warren had participants listen to a recording of the sentence "The state governors met with their respective legislatures convening in the capital city." Warren replaced the first /s/ in "legislatures" with the sound of a cough and told his participants that they should indicate where in the sentence the cough occurred. No participant identified the correct position of the cough, and, even more significantly, none of them noticed that the /s/ in "legislatures" was missing. This effect, which Warren called the **phonemic restoration effect**, was experienced even by students and staff in the psychology department who knew that the /s/ was missing. Participants "filled in" the missing phoneme based on the context produced by the sentence and the portion of the word that was presented.

Warren also showed that the phonemic restoration effect can be influenced by the meaning of the words that *follow* the missing phoneme. For example, the last word of the phrase "There was time to *ave . . ." (where the * indicates the presence of a cough or some other sound) could be *shave*, *save*, *wave*, or *rave*, but participants heard the word *wave* when the remainder of the sentence had to do with saying good-bye to a departing friend. Thus, our perception of speech is influenced by top-down processing—our knowledge of the meanings of words that we bring to the situation. The effect of top-down processing has also been demonstrated by finding that more restoration occurs for a real word like prOgress (where the capital letter indicates the masked phoneme) than for a similar "pseudoword" like crOgress (Samuel, 1990). We will now consider another example of how our knowledge of the meanings of words helps us to perceive them.

Speech Segmentation The words on this page are easy to recognize. Each word is separated by a space, so it's easy to tell one word from another. However, when people hear words in a conversation, these words are not separated by spaces, or pauses, even though it may sound like they are.

When we look at a record of the physical energy produced by conversational speech, we see that the speech signal is continuous, with either no physical breaks in the signal or breaks that don't correspond to the breaks we perceive between words (Figure 10.2). The fact that there are usually no spaces between words becomes obvious when you listen to someone speaking a foreign language. To someone who is unfamiliar with that language, the words seem to speed by in an unbroken string. However, to a speaker of

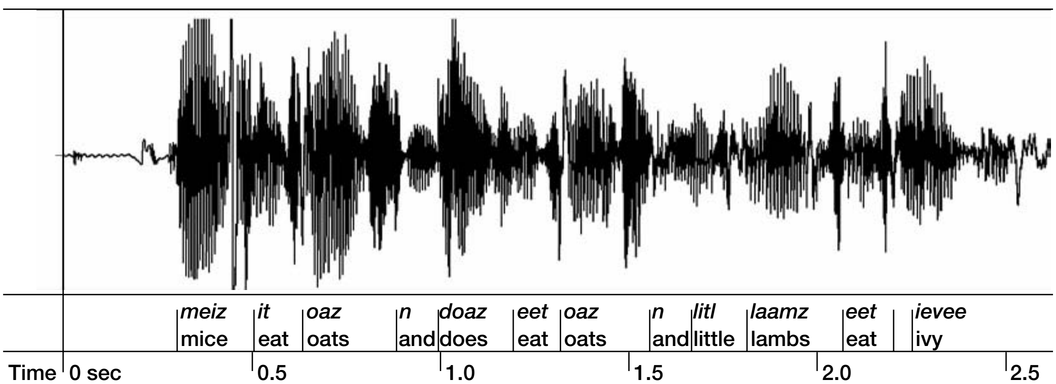

■ **Figure 10.2** Sound energy for the phrase "Mice eat oats and does eat oats and little lambs eat ivy." The italicized words just below the sound record indicate how this phrase was pronounced by the speaker. The vertical lines next to the words indicate where each word begins. Note that it is difficult or impossible to tell from the sound record where one word ends and the other begins. (Speech signal courtesy of Peter Howell.)

that language, the words seem separated, just as the words of languages you know seem separated to you. The process of perceiving individual words from the continuous flow of the speech signal is called **speech segmentation** (see Chapter 3, page 82).

Our ability to achieve speech segmentation is made more complex by the fact that not everyone produces words in the same way. People talk with different accents and at different speeds, and most important, people often take a relaxed approach to pronouncing words when they are speaking naturally. For example, how would you say "Did you go to class today," if you were talking to a friend? Would you say "Did you" or "Dijoo"? You have your own ways of producing various words and phonemes, and other people have theirs. For example, analysis of how people actually speak has determined that there are 50 different ways to pronounce the word *the* (Waldrop, 1988).

The way people pronounce words in conversational speech makes about half of the words unintelligible when taken from their fluent context and presented alone. Irwin Pollack and J. M. Pickett (1964) demonstrated this by recording the conversations of participants who sat in a room, waiting for the experiment to begin. When the participants were then presented with recordings of single words taken out of their own conversations, they could identify only half the words, even though they were listening to their own voices!

There are a number of types of information that listeners can use to deal with the problems posed by words in spoken sentences. One of these is the context, or the meaning, of a conversation. The importance of context is illustrated by the results of the Pollack and Pickett experiment, because it showed that when words are taken out of the context provided by other words in a conversation, understanding the words becomes much more difficult.

Our understanding of meaning also helps solve the problem of speech segmentation. An unfamiliar language that sounds like an unbroken string of sounds becomes segmented into individual words once you learn the language. When you learn the language, you not only learn meanings but you also learn that certain sounds are more likely to occur at the ends or beginnings of words. For example, in English, words can end in *rk* (*work*, *fork*), but not *kr*. However, words can begin with *kr* (*krypton*, *krill*), but not *rk*. There is evidence that people learn these rules about permissible beginnings and endings of words as young children (Gomez & Gerkin, 1999, 2000; Saffran et al., 1999). As we saw in Chapter 3, infants as young as 8 months of age can achieve speech segmentation through a process called *statistical learning*. We now move from *perceiving* letters and words to factors that influence our ability to *understand* words.

Understanding Words

Our ability to understand words is influenced by a number of factors, including how common the word is and the other words that surround it in a sentence.

Word Frequency

An adult's lexicon may contain over 50,000 words, but some of these words can be more easily accessed than others. One factor that contributes to these differences in accessibility is **word frequency**—the relative usage of a word in a particular language. For example, in English, *home* occurs 547 times per million words, and *hike* occurs only 4 times per million words. The **word-frequency effect** refers to the fact that we respond more rapidly to high-frequency words like *home* than to low-frequency words like *hike*. One way this has been demonstrated is through the *lexical decision task* (see Method: Lexical Decision Task, page 303).

Demonstration

Lexical Decision Task

CogLab

Lexical Decision

In the **lexical decision task**, a participant reads a list that consists of words and nonwords. Your task is to indicate as quickly as possible whether each entry in the lists below is a word. Try this yourself by silently reading List 1 below and saying "yes" each time you encounter a word. Either time yourself to determine how long it takes to get through the list or just notice how difficult the task is.

List 1

Gambastya, revery, voitle, chard, wefe, cratily, decoy, puldow, faflot, oriole, voluble, boovle, chalt, awry, signet, trave, crock, cryptic, ewe, himpola.

Now try the same thing for List 2:

List 2

Mulvow, governor, bless, tuglety, gare, relief, ruftily, history, pindle, develop, grdot, norve, busy, effort, garvola, match, sard, pleasant, coin, maisle.

The task you have just completed (which is taken from D. W. Carroll, 1999; also see Hirsh-Pasek et al., 1993) is called a lexical decision task because you had to decide whether each group of letters was a word in your lexicon.

When researchers presented this task under controlled conditions, they found that people read high-frequency words faster than low-frequency words (Savin, 1963). Thus, it is likely that you were able to carry out the lexical decision task more rapidly for List 2 compared to List 1.

This slower response for less-frequent words has also been demonstrated by measuring people's eye movements as they are reading (see Method: Measuring Eye Movements, page 120). The eye movements that occur during reading consist of fixations, during which the eye stops on a word for about a quarter of a second (250 ms), and movements, which propel the eye to the next fixation.

In a recent eye-movement study, Keith Rayner and coworkers (2003) had participants read sentences that contained either a high- or a low-frequency target word. For example, the sentence "Sam wore the horrid coat though his pretty girlfriend complained," contains the high-frequency target word *pretty*. The other version of the sentence was exactly the same, but with the high-frequency word *pretty* replaced by the

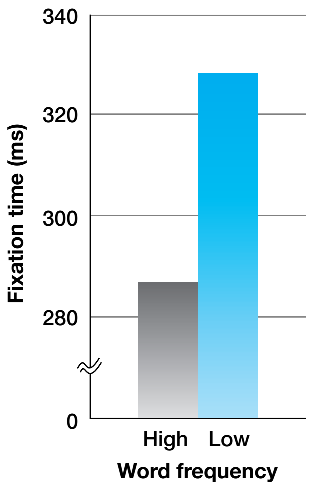

Figure 10.3 Results of Rayner et al.'s (2003) experiment. The bars indicate how long participants looked at target words like *pretty* and *demure.* These results show that participants fixated low-frequency words longer than high-frequency words.

low-frequency word *demure*. The results, shown in Figure 10.3, indicate that readers looked at the low-frequency words about 40 ms longer than the high-frequency words.

Context Effects Our ability to access words in a sentence is affected not only by frequency, but also by the meaning of the rest of the sentence. As we will see when we consider how we understand sentences, we are constantly attempting to figure out what a sentence means as we are reading it. This process involves both understanding individual words and understanding how these words fit into the overall meaning of the sentence. For example, it takes less time to understand *The Eskimos were frightened by the walrus* than to understand *The bankers were frightened by the walrus*, because words that are expected within the context of the sentence (like *walrus* appearing with *Eskimos*) are understood more rapidly than words that are not expected (like *walrus* appearing with *bankers*; Marslen-Wilson, 1990).

Lexical Ambiguity Words can often have more than one meaning, a situation called **lexical ambiguity**. For example, the word *bug* can refer to insects, or hidden listening devices, or being annoyed, among other things. When ambiguous words appear in a sentence, we usually use the context of the sentence to determine which definition applies. For example, if Susan says "My mother is bugging me," we can be pretty sure that *bugging* refers to the fact that Susan's mother is annoying her, as opposed to sprinkling insects on her or installing a hidden listening device in her room (although we might need further context to totally rule out this last possibility).

Context often clears up ambiguity so rapidly that we are not aware of its existence. However, David Swinney (1979) showed that people briefly access multiple meanings of ambiguous words before the effect of context takes over. He did this by presenting participants with a tape recording of sentences such as the following:

> 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 the room.

If you had to predict which meaning listeners would use for *bugs* in this sentence, *insect* would be the logical choice because the sentence mentions spiders and roaches. However, using a technique called *lexical priming*, Swinney found that right after the word *bug* was presented, his listeners had accessed two meanings.

Method

Lexical Priming

Remember from Chapter 5 that priming occurs when seeing a stimulus makes it easier to respond to that stimulus when it is presented again (See Method: Repetition Priming, page 192). The basic principle behind priming is that the first presentation of a stimulus activates a representation of the stimulus, and a person can respond more rapidly to the stimulus if this activation is still present when the stimulus is presented again.

Priming involving the naming of words is called **lexical priming**. Because lexical priming involves the *meaning* of words, priming effects can occur when a word is followed by another word with a similar meaning. For example, presenting the word *ant* before presenting the word *bug* can cause a person to respond faster to the word *bug* than if *ant* had not preceded it. The presence of a lexical priming effect would, therefore, indicate whether two words, like *ant* and *bug,* have similar meanings in a person's mind.

Swinney used lexical priming by presenting the passage about the government building to participants and, as they were hearing the word *bug*, presenting a word or a nonword on a screen (Figure 10.4a). The words he presented were either related to the "insect" meaning of *bug* (*ant*), or to the "hidden listening device" meaning (*spy*), or were not related at all (*sky*). The participant was told to indicate as quickly as possible whether the item flashed on the screen was a word or a nonword. (See Method: Lexical Decision Task, page 364.)

Swinney's result, shown in Figure 10.4b, was that participants responded with nearly the same speed to both *ant* and *spy* (the small difference between them is not significant), and the response to both of these words was significantly faster than the response to *sky*. This faster responding to words associated with two of the meanings of *bug* means that even though there is information in the sentence indicating that *bug* is an insect, listeners accessed both meanings of *bug* as it was being presented. This effect was, however, short-lived, because when Swinney repeated the same test but waited for two or three

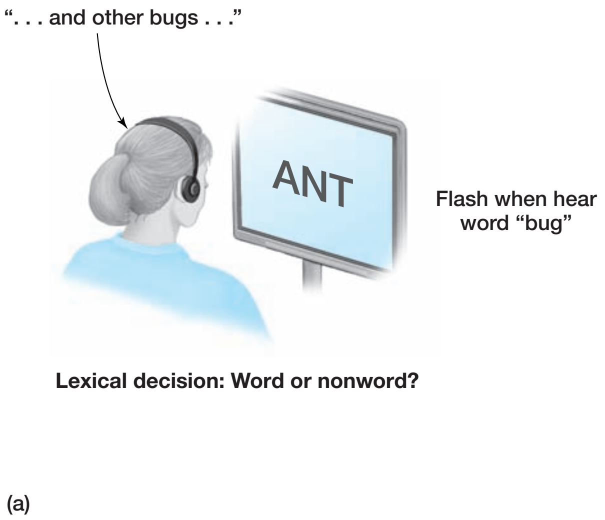

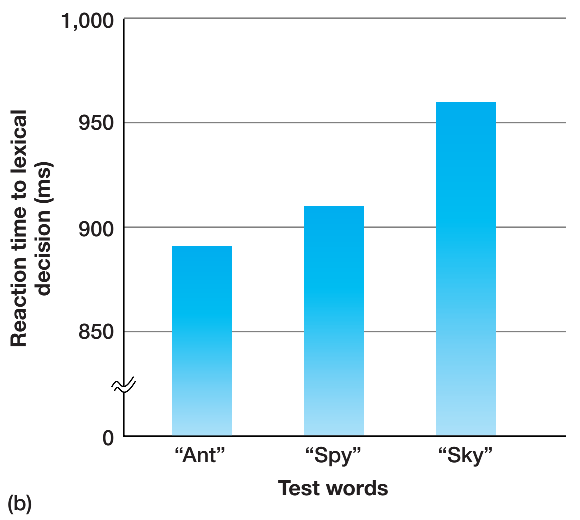

Figure 10.4 (a) The procedure for Swinney's (1979) experiment. See text for details. (b) The results of Swinney's experiment. The fact that the reaction times to *ant* and *spy* were not significantly different showed that people briefly accessed both meanings of the word *bugs* as they read this word in a sentence.

syllables before presenting the test words, the effect had vanished. Thus, within about 200 ms after hearing *bug*, the *insect* meaning had been selected from the ones initially activated. Context does, therefore, have an effect on the activation of word meaning, but it exerts its influence after all meanings of a word have been briefly accessed.

Summary: Words Alone and in Sentences

Figure 10.5 summarizes the results we have described for perceiving letters and words, and Figure 10.6 summarizes the results for accessing words. Note that for all of the effects we discussed (except for the word-frequency effect), the meanings of words fa-

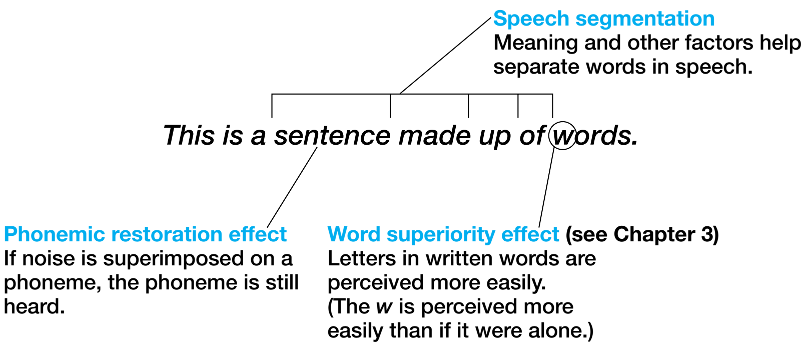

■ **Figure 10.5** Summary of the two effects we described that influence the perception of letters and words: (1) speech segmentation, and (2) the phonemic restoration effect. The word-superiority effect from Chapter 3 (see p. 64) is also included.

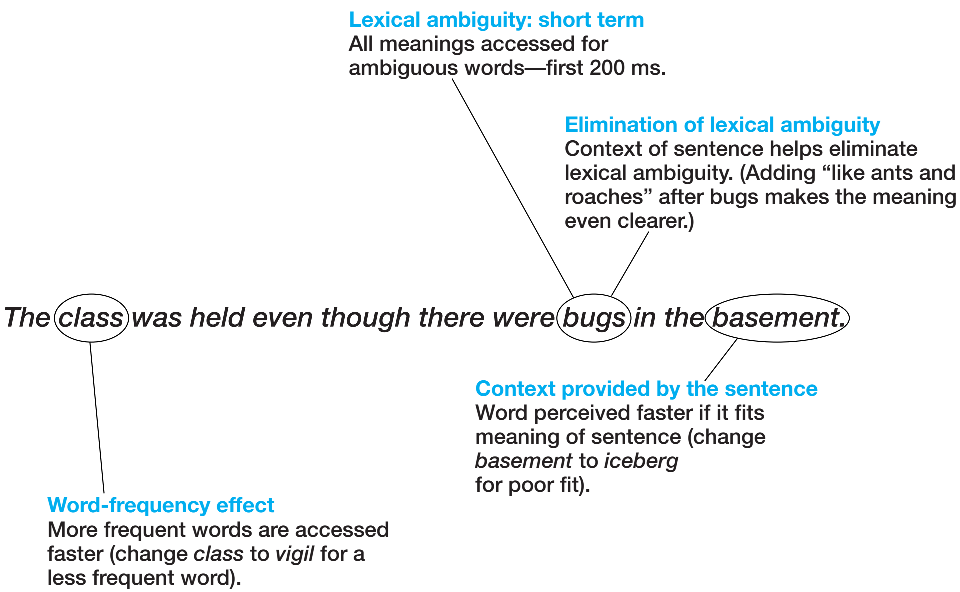

■ **Figure 10.6** Summary of the four effects we described in connection with accessing words: (1) short-term lexical ambiguity, (2) elimination of lexical ambiguity, (3) how the context of a sentence can cause words to be perceived faster, and (4) the word-frequency effect.

cilitated perceiving letters and phonemes, and the meaning of sentences facilitated understanding words. These results are important because they illustrate one of the main messages of this chapter: Although the study of language is often described in terms of its individual components—such as letters, words, and sentences—these components are not processed in isolation. As we discuss how we understand sentences, we will see more examples of how each of these components interacts with and influences one another.

Understanding Sentences

Although the last section was about words, we ended up discussing sentences as well. This isn't surprising because words rarely appear in isolation. They appear together in sentences, with all of the words combining to create the meaning of the sentence. To understand how words work together to create the meaning of the sentence, we first need to distinguish between two properties of sentences: semantics and syntax.

Semantics and Syntax

Semantics is the meanings of words and sentences. **Syntax** is the rules for combining words into sentences. Recent experiments have demonstrated a physiological distinction between these two characteristics of words and sentences. For example, semantics and syntax are associated with different components of a physiological response called the *event-related potential (ERP).*

Method

Event-Related Potential

Most stimuli activate many thousands of neurons in the brain. The signals generated by these neurons can be measured in humans by recording the event-related potential with disk electrodes placed on a person's scalp (Figure 10.7a). When a stimulus is presented, the electrodes record voltage changes in the brain that are generated by the thousands of neurons near each electrode.

One thing that makes the ERP a valuable tool for cognitive psychology is that the response consists of a number of different components, which occur at different delays after a stimulus is presented. Figure 10.7b shows the N400 component of the ERP. "N" stands for "negative" (note that negative is up in ERP records), and 400 stands for the time at which the response peaks—400 ms from the presentation of the stimulus in this case. The N400 component is influenced by whether a word fits the meaning of a sentence. For example, the colored line in Figure 10.7b shows the N400 response to the word "eat" in the sentence "The cats won't eat." The gray line shows the response to the word "bake" in "The cats won't

Image not available due to copyright restrictions

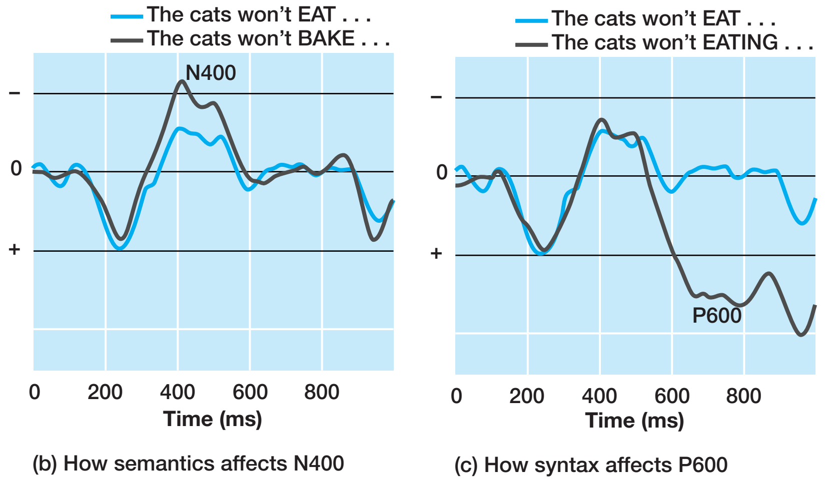

Figure 10.7 (b) The N400 wave of the ERP is affected by semantics. It becomes larger (dark gray line) when the meaning of a word does not fit the rest of the sentence. (c) The P600 wave of the ERP is affected by syntax. It becomes larger (dark gray line) when syntax is incorrect. (Parts b and c: Reprinted from *Trends in Cognitive Sciences,* Volume 1, Issue 6, Osterhout et al., "Event-Related Potentials and Language," Figure 1, Copyright © 1997 with permission from Elsevier.)

bake." The N400 response increases to "bake" because the word "bake" doesn't fit in this sentence (Neville et al., 1991; Osterhout et al., 1997; also see Kutas & Federmeier, 2000).

The fact that the N400 response is sensitive to the meaning of a word in a sentence means that this response is associated with semantics. Figure 10.7c shows that the P600 wave is associated with violations of syntax. This wave is small when syntax is correct, but becomes larger when syntax is incorrect. Thus, "The cats won't eating . . ." causes a larger P600 response than "The cats won't eat. . . ." The fact that semantics and syntax are associated with different waves of the ERP supports the idea that they are associated with different mechanisms.

The idea that semantics and syntax are associated with different mechanisms has also been supported by brain imaging studies, which have shown that different areas of the brain are activated by semantics and syntax (Dapretto & Bookheimer, 1999). Also, damage to some areas of the brain causes difficulties in understanding the meanings of words, and damage to other areas causes problems understanding grammar (Breedin & Saffran, 1999).

We will see that semantics and syntax interact with one another as a reader or listener works to determine the meaning of a sentence. One of the central processes for determining meaning is **parsing**, the mental grouping of words in a sentence into phrases.

Parsing a Sentence

The goal of parsing is to determine the message of a sentence. This message is determined by the meanings of the words in the sentence and how these words are grouped together into phrases. As we will see, a number of groupings are often possible for the same string of words. Consider for example, the sentence

> The spy saw the man with the binoculars.

As people begin reading a sentence such as this one, they make guesses about how the sentence is going to unfold. Upon reading "The spy saw the man . . . ," it seems as if there is only one meaning—a spy looking at a man. However, as the sentence continues, the phrase ". . . with the binoculars" poses a problem, because now there are two possible meanings:

> *Meaning 1:* The spy was looking through some binoculars to see the man. This meaning groups "the spy" and "the binoculars," like this:

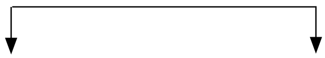

> *The spy* saw the man *with the binoculars.*

or

> *Meaning 2:* The spy was looking at a man, who had some binoculars. This meaning groups "the man" and "the binoculars," like this:

> The spy saw *the man with the binoculars.*

This sentence provides an example of **syntactic ambiguity**—the words are the same, but there is more than one possible structure, and so there is more than one meaning. Although there is no way to know the correct meaning of this sentence from the information given, there is a tendency for people to interpret this sentence in terms of the first meaning, with the spy being the one with the binoculars.

What causes us to prefer one way of parsing the sentence over another? Psychologists have proposed that there is a mechanism responsible for determining the meaning of the sentence. This mechanism has been called a number of things, including the *language-analysis device* and the *sentence-analyzing mechanism.* We will simply call it the **parser**. The parser determines the meaning of the sentence, primarily by determining how words are grouped together into phrases. Psychologists are interested in answering the question: "What factors determine how the parsing mechanism works?" Two answers have been proposed to this question; one assigns the central role to syntax, with semantics coming into play later, and the other proposes that syntax and semantics work simultaneously to determine the meaning of a sentence.

The Syntax-First Approach to Parsing As its name implies, the **syntax-first approach to parsing** focuses on how parsing is determined by syntax—the grammatical structure of the sentence. We can appreciate a connection between syntax and sentence understanding by considering the following poem from Lewis Carroll's (1872) *Through the Looking Glass and What Alice Found There:*

'Twas brillig, and the slithy toves
Did gyre and gimble in the wabe:
All mimsy were the borogoves,
And the mome raths outgrabe.

Even though the words in this poem are nonsense, we still have a sense of what the poem is about. The first two lines seem to be about creatures called *slithy toves* doing something called *gyring and gimbling* in a place called *the wabe*. We are able to create meaning out of gibberish because we use syntax to infer meaning (Kako & Wagner, 2001).

The syntax-first approach to parsing states that the parsing mechanism groups phrases together based on structural principles. One of these principles is called *late closure*. The principle of **late closure** states that when a person encounters a new word, the parser assumes that this word is part of the current phrase (Frazier, 1987). We can illustrate this by considering the following sentence:

Because he always jogs a mile seems like a short distance to him.

Table 10.1 indicates how you may have read this sentence. At first, this sentence seems to be about a man who jogs a mile (a) and (b), but trouble occurs when you get to *seems*

Table 10.1 The Principle of Late Closure

First Try	
Part of the Sentence	**Probable Reader's Reaction**
(a) Because he always jogs	This is about a man who jogs.
(b) a mile	He jogs a mile.
(c) seems like	This doesn't make sense. How does "seem like" fit in here?
(d) a short distance to him.	OK. I read the sentence incorrectly the first time. I'll try again.
Second Try	
Part of the Sentence	**Probable Reader's Reaction**
Because he always jogs	The man jogs.
a mile seems like a short distance to him.	He is in good shape so a mile doesn't seem like much.

like (c) and after reading *a short distance to him* (d), you realize that there is another way to read the sentence (see "Second Try," Table 10.1). Because this sentence has led the reader "down the garden path" (down a path that seems right, but turns out to be wrong), this sentence is called a **garden-path sentence** (Frazier & Rayner, 1982).

We can see how this sentence illustrates the principle of late closure by focusing on the words *a mile*. According to the late-closure principle, the parser assumes that *a mile* is a continuation of the phrase *because he always jogs*. However, in reality, *a mile* is the beginning of a new phrase. Late closure (so named because it proposes to keep adding new words to the current phrase, so it delays closing off the phrase for as long as possible) leads to the wrong parse—the phrase needed to be closed after *jogs*, so the new phrase can begin. (Note that for this sentence, the correct phrasing could be indicated by inserting a comma after *jogs*. However, there are other garden-path sentences, such as "The student knew the answer to the question was wrong," that cannot be made easier to understand by adding a comma.) Because application of the syntactic rule of late closure results in a garden-path sentence, the syntax-first approach to parsing has also been called the **garden-path model** (Frazier & Rayner, 1982).

A number of experiments show that parsing is determined by late closure and other syntactic principles (Frazier, 1987). Although the garden-path model of parsing focuses on syntax, it doesn't ignore semantics. It states that if analysis in terms of syntax doesn't make sense, then semantics is used to clear up the ambiguity. Thus, according to this approach to parsing, syntax is used first, then semantics is called on, if needed, to make sense of the sentence (as it did in the "Second Try" in Table 10.1).

We can draw a comparison between this process of determining how words in a sentence are grouped into phrases and how parts of a visual scene become perceptually grouped into objects. Remember our example from Chapter 3 of how the Gestalt principles of organization help us guess that the scene shown in Figure 10.8a might be a

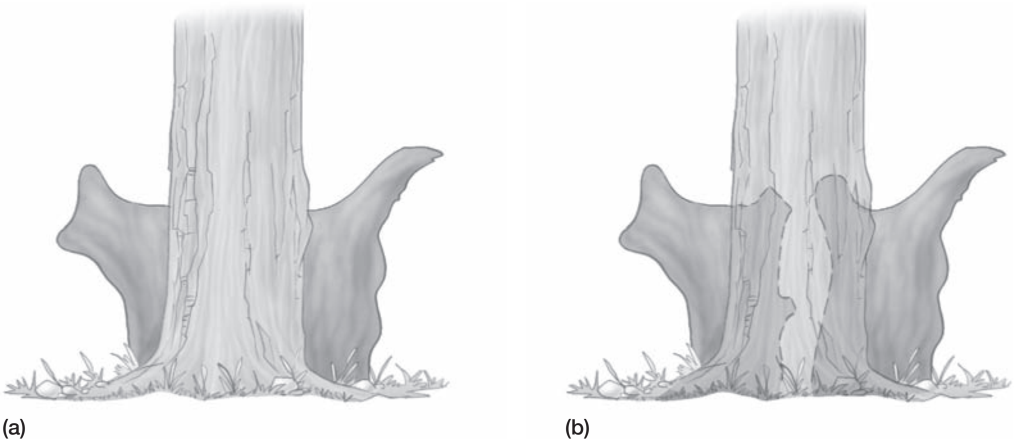

Figure 10.8 (a) What lurks behind this tree? (b) It is *not* a creature!

creature hiding behind a tree. As it turns out, further information provided by looking behind the tree proves that guess to be wrong and so we revise our assessment of the situation from "creature hiding behind a tree" to "strange tree trunks behind a tree" (Figure 10.8b).

We used this example in the perception chapter to illustrate the idea that the Gestalt laws of organization are heuristics—rules of thumb that are "best-guess" rules for determining our perceptions. Most of the time, these rules result in perceptions that provide accurate information about what is "out there," and they have the advantage of being fast, which is essential because our very survival depends on quickly reacting to objects and events in the environment.

A similar process occurs when our language system uses a rule such as the principle of late closure to provide a "best guess" about the unfolding meaning of a sentence. Most of the time, this rule leads to the correct conclusion about how the sentence should be parsed. However, in some cases, such as when psychologists create garden-path sentences like the one about the jogger, the rule results in an incorrect parsing, which has to be corrected when more information becomes available at the end of the sentence.

Thus, just as ambiguous visual scenes help perception researchers understand the processes involved in visual perception, garden-path sentences help language researchers determine the processes involved in understanding language. Garden-path sentences accomplish this by showing us what guesses the parser makes, as in the sentence about the jogger (Fodor, 1995).

The Interactionist Approach to Parsing An indication of the role of semantics in understanding sentences is provided by comparing "The spy saw the man with the binoculars" to "The bird saw the man with the binoculars." We have seen that the sentence about the spy has two meanings: The spy could be looking at a man through the binoculars or could be looking at a man who has a pair of binoculars. However, by changing *spy* to *bird*, we create a sentence with only one reasonable meaning because we know the bird wouldn't be looking at the man through binoculars. Thus, in this sentence, our knowledge of birds makes it clear that the man is the one with the binoculars, and not the bird.

According to the **interactionist approach to parsing**, semantics can influence processing as the person is reading the sentence. *All* information, *both syntactic and semantic*, is taken into account as we read a sentence, so any corrections that need to occur in the processing of a sentence take place as the person is reading the sentence (Altmann, 1998; Altmann & Steedman, 1988; MacDonald et al., 1994). Thus, the crucial question in comparing the syntax-first and interactionist approaches is not *whether* semantics is involved, but *when* semantics comes into play. Is semantics activated only after syntax has determined the initial parsing (syntax-first approach), or does semantics come into play as a sentence is being read (interactionist approach)?

Recently, a number of studies in which readers' or listeners' eye movements have been measured while they are reading or listening to a sentence have helped answer this question (see Method: Eye Movements, page 120). In one study, Michael Tanenhaus and coworkers (1995) used eye movements to study how people process the information in sentences. They presented a picture that illustrated the objects mentioned in the sentence and then determined where participants looked while they were trying to understand the sentence. One of the sentences was

Put the apple on the towel in the box.

The beginning of this sentence (*Put the apple on the towel*) sounds like a straightforward request to put an apple on a towel. But after hearing the last part of the sentence (*in the box*), two possible meanings emerge: The sentence could be about *where* to put the apple (put it on the towel that's inside the box; Figure 10.9a), or about *which* apple (pick the apple that is on the towel to put in the box; Figure 10.9b).

Tanenhaus reasoned that in most real-life situations we hear sentences while we are interacting with the environment. Thus, the purpose of this experiment was to see how the environmental context that accompanies a sentence can influence how a person moves their eyes to fixate on particular objects in the environment. Tanenhaus used two pictures. The picture in Figure 10.10a is the "two-apple condition," in which one apple is shown on a towel and the other is on a napkin. Figure 10.10b is the "one-apple condition," in which one apple is shown on a towel.

The eye movements for participants looking at the two-apple condition, indicated by the arrows, shows that when participants heard "Put the apple," they moved their eyes to the apple on the napkin (arrow A). When they heard "on the towel," they moved their eyes to the apple on the towel (B), and when they heard "in the box," they looked at the box (C).

Participants in the one-apple condition responded differently. After hearing "Put the apple," they moved their eyes to the apple (A). After "on the towel," they looked at

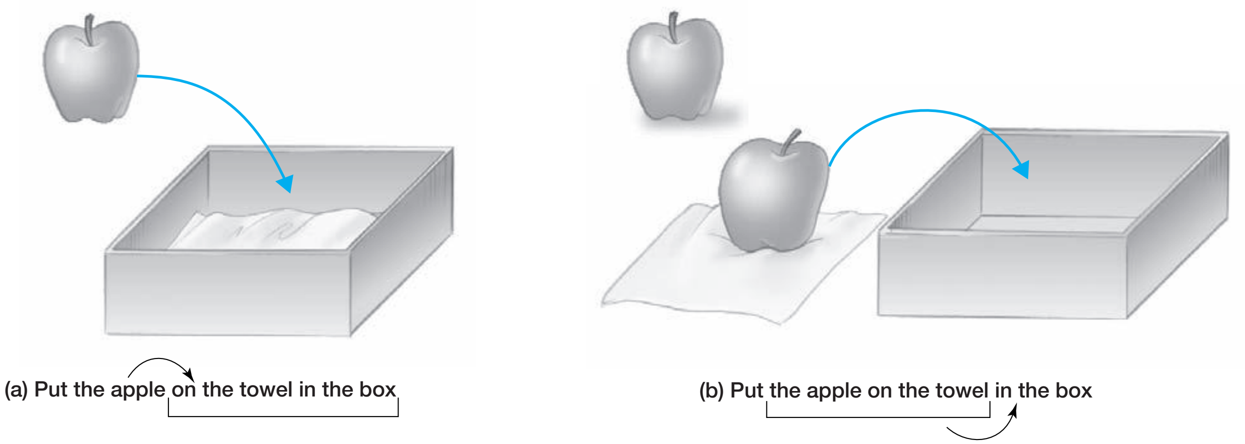

■ **Figure 10.9** These two pictures indicate two meanings of the sentence "Put the apple on the towel in the box," which was used in Tanenhaus et al.'s (1995) eye-movement study. (a) The apple goes on the towel that's inside the box. (b) The apple on the towel goes in the box.

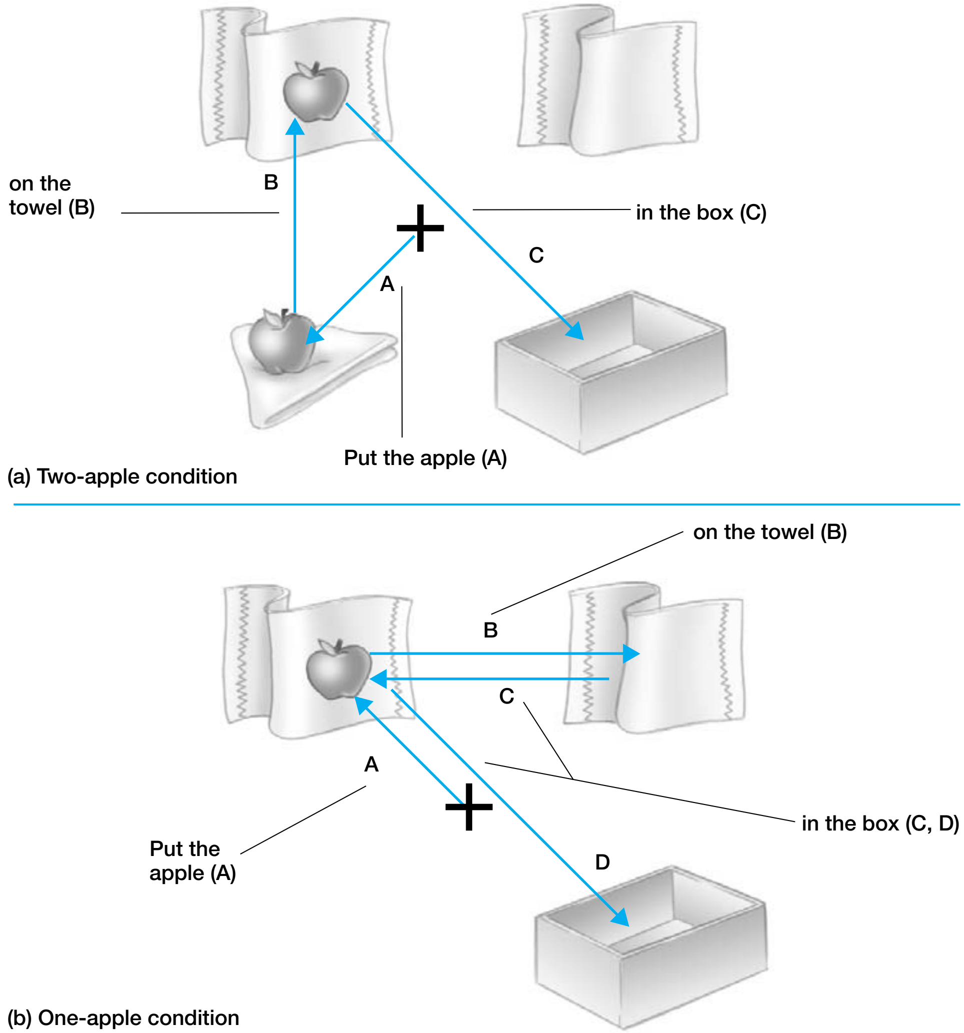

■ **Figure 10.10** The results of Tanenhaus et al.'s (1995) eye-movement study. The way participants moved their eyes to different parts of the pictures depended on the information provided by the picture. Note how the eye movements differ in (a) the two-apple condition and (b) the one-apple condition.

the other towel (B). However, upon hearing "in the box," they apparently realized that the sentence was asking them to move the apple not to the other towel, but to the box, so they quickly moved their eyes back to the apple (C) and then to the box (D).

In this experiment, the participants' eyes movements provided information about what they were thinking. We can summarize the two conditions as follows: When there were two apples, participants initially thought "*On the towel* means I should pick the apple that is on the towel." (They are thinking *which* apple to pick). But when there was just one apple, participants initially thought "*On the towel* means I should place the apple on the towel." (They are thinking *where* to place the apple). The important result of this experiment is that in the one-apple condition, participants' eye movements changed as soon as they received information that indicated that they needed to revise their initial interpretation of the sentence. This immediate responding supports the interactionist idea that the reader or listener takes both syntactic and semantic information into account simultaneously. (Also see Altmann & Kamide, 1999, for another demonstration of how the eyes rapidly respond to the meaning of a sentence.)

Although the controversy regarding whether the syntax-first approach or the interactionist approach is correct is still not resolved (Bever et al., 1998; Rayner & Clifton, 2002), evidence such as the results of the eye-movement study indicates that semantics can be taken into account earlier than proposed by the syntax-first approach. Furthermore, the result of the "apple in the box" experiment also shows that information in addition to the words in the sentence help determine what a sentence means. This is important, because in real life we rarely encounter sentences in isolation. Rather, we encounter sentences while we are in specific environments, or as we are listening to conversations or reading a story.

That sentences occur within a context is particularly important for reading, because sentences are typically part of a larger text or story. Thus, when we read a particular sentence, we already know a great deal of information about what is happening from what we read before. This brings us to the next level of the study of language—the study of how we understand text and stories (commonly called *discourse processing* or *text processing*). As we will see, most research in text processing is concerned with how readers' understanding of a story is determined by information provided by many sentences taken together.

Test Yourself 10.1

1. What is special about human language? Consider why human language is unique and what it is used for.
2. What events are associated with the beginning of the modern study of language in the 1950s?
3. What is psycholinguistics? What are its concerns, and what part of psycholinguistics does this chapter focus on?
4. What evidence supports the statement that "meaning makes it easier to perceive letters in words, and words in spoken sentences"?
5. How do the frequency of words and the context of a sentence aid in accessing words? How does Swinney's experiment about "bugs" indicate that the meanings of ambiguous words can take precedence over context, at least for a short time?
6. Why do we say that there is more to understanding sentences than simply adding up the meanings of the words that make up the sentence?

7. Describe the syntax-first explanation of parsing and the interactionist explanation. What are the roles of syntax and semantics in each explanation? What evidence supports the interactionist approach?

Understanding Text and Stories

Just as sentences are more than the sum of the meanings of individual words, stories are more than the sum of the meanings of individual sentences. In a well-written story, sentences in one part of the story are related to sentences in other parts of the story. Thus, the reader's task is to use these relationships between sentences to create a coherent, understandable story.

The materials used in research on text processing are usually excerpts from narrative texts. *Narrative* refers to texts in which there is a story that progresses from one event to another, although stories can also include flashbacks of events that happened earlier. An important property of any narrative is **coherence**—the representation of the text in a person's mind so that information in one part of the text is related to information in another part of the text. Texts with coherence are usually easier to understand than texts without coherence.

How Inference Creates Coherence

Most of the coherence in texts is created by inference. **Inference** refers to the process by which readers create information during reading that is not explicitly stated in the text. We have had a great deal of experience with inference from our study of memory in Chapter 7. For example, on page 255 we described an experiment in which participants read the passage "John was trying to fix the bird house. He was pounding the nail when his father came out to watch him do the work." We saw that after reading that passage, participants were likely to say that they had previously seen the following passage: "John was using a hammer to fix the birdhouse. He was looking for the nail when his father came out to watch him." The fact that they thought they had seen this passage, even though they had never read that John was using a hammer, occurred because they had inferred that John was using a hammer from the information that he was *pounding the nail* (Bransford & Johnson, 1973). People use a similar creative process to make a number of different types of inferences as they are reading a text.

Anaphoric Inference

Inferences that connect an object or person in one sentence to an object or person in another sentence are called **anaphoric inferences**. For example, consider the following.

> Riffifi, the famous poodle, won the dog show. She has now won the last three shows she has entered.

Anaphoric inference occurs when we infer that *She* at the beginning of the second sentence and the other *she* near the end both refer to Riffifi. In the previous "John and the birdhouse" example, knowing that *He* in the second sentence refers to *John* is another example of anaphoric inference.

We usually have little trouble making anaphoric inferences because of the way information is presented in sentences and our ability to make use of knowledge we bring to the situation. But here is an example of a quote from a *New York Times* interview with former heavyweight champion George Foreman (who has recently been known for lending his name to a popular line of grills), which puts our ability to create anaphoric inference to the test.

> What we really love to do on our vacation time is go down to our ranch in Marshall, Texas. We have close to 500 acres. There are lots of ponds and I take the kids out and we fish. And then, of course, we grill them. (Stevens, 2000)

Based just on the structure of the sentence, we might conclude that the kids were grilled, but we know that the chances are pretty good that the fish were grilled, not George Foreman's children! Readers are capable of creating anaphoric inferences even under adverse conditions because they add information from their knowledge of the world to the information provided in the text.

Instrumental Inference Inferences about tools or methods are **instrumental inferences**. For example, when we read the sentence "William Shakespeare wrote Hamlet while he was sitting at his desk," we infer from what we know about the time during which Shakespeare lived that he was probably using a quill pen (not a laptop computer!) and that his desk was made of wood. Similarly, inferring from the passage about John that he is using a hammer to pound the nails would be an instrumental inference.

Causal Inference Inferences that result in the conclusion that the events described in one clause or sentence were caused by events that occurred in a previous sentence are **causal inferences** (Goldman et al., 1999; Graesser et al., 1994; van den Broek, 1994). For example, when we read the sentences

> Sharon took an aspirin. Her headache went away.

we infer that the aspirin caused the headache to go away (Singer et al., 1992). This is an example of a fairly obvious inference that most people in our culture would make based on their knowledge about headaches and aspirin.

Other causal inferences are not so obvious and may be more difficult to figure out. For example, what do you conclude from reading the following sentences?

> Sharon took a shower. Her headache went away.

You might conclude, from the fact that the headache sentence directly follows the shower sentence, that the shower had something to do with eliminating Sharon's headache. However, the causal connection between the shower and the headache is weaker than the connection between the aspirin and the headache in the first pair of sentences. Making the shower–headache connection requires more work for the reader. You might infer that the shower relaxed Sharon, or perhaps her habit of singing in the shower was therapeutic. Or you might decide there actually isn't much of a connection between the two sentences.

Causal inferences create connections that are essential for creating coherence in texts, and making these inferences can involve creativity by the reader. Thus, reading a text involves more than just understanding words or sentences. It is a dynamic process that involves transformation of the words, sentences, and sequences of sentences into a meaningful story (Goldman et al., 1999; Graesser et al., 1994; van den Broek, 1994). Sometimes this is easy, sometimes harder, depending on the skill and intention of the both the reader and writer.

We have, so far, been describing the process of text comprehension in terms of connecting the meanings of sentences to create a story. Another approach to understanding text comprehension is to look directly at the nature of the mental representation that people form as they read a story. This is called the *situation model* approach to text comprehension.

Situation Models

A **situation model** is a mental representation of what a text is about. This approach proposes that the mental representation people form as they read a story does not indicate information about phrases, sentences, or paragraphs, but, instead, includes a representation of the situation in terms of the people, objects, locations, and events that are being described in the story (Graesser & Wiemer-Hastings, 1999; Zwaan, 1999). The situation-model approach also proposes that readers vicariously experience events that are being described in a story and that this experience is often from the point of view of the protagonist—the main character in the story or the character being described at a particular point in the story.

For example, in a story about a man walking through a building, the reader would create a map of the space through which the protagonist is walking and keep track of the protagonist's location in the building. According to the idea of situation models, if specific objects are described in the story, then the reader will have better access to information about objects that are near to the protagonist or are more visible to the protagonist (Morrow et al., 1987).

This way of looking at how readers process stories predicts that information about objects or events that are difficult for the protagonist to access will also be difficult for the reader to access. An experiment by William Horton and David Rapp (2003) tested this idea using short passages like the following:

1. Melanie ran downstairs and threw herself onto the couch.
2. An exciting horror movie was on television.
3. She opened a bag of chips and dug right in.
4. She watched a vampire stalk the helpless victim.
5. She had never seen this movie before.

Participants are then presented with one of the following endings:

Blocked story continuation (Figure 10.11a):

6a. Melanie's mother appeared in front of the TV.
7a. She told Melanie not to forget about her homework.

or

Unblocked story continuation (Figure 10.11b):

6b. Melanie's mother appeared behind the TV.
7b. She told Melanie not to forget about her homework.

Participants read the story line by line from a computer screen. After sentence 7, a warning tone sounded, which indicated that the target question was going to be presented. The target question for the story above was "Was the victim being stalked by a vampire?" The participant's task was to answer "yes" or "no" as quickly as possible by pressing the correct key on the computer keyboard.

The situation-model prediction is that participants who read the blocked story continuation should react more slowly to the test question because the TV screen, which

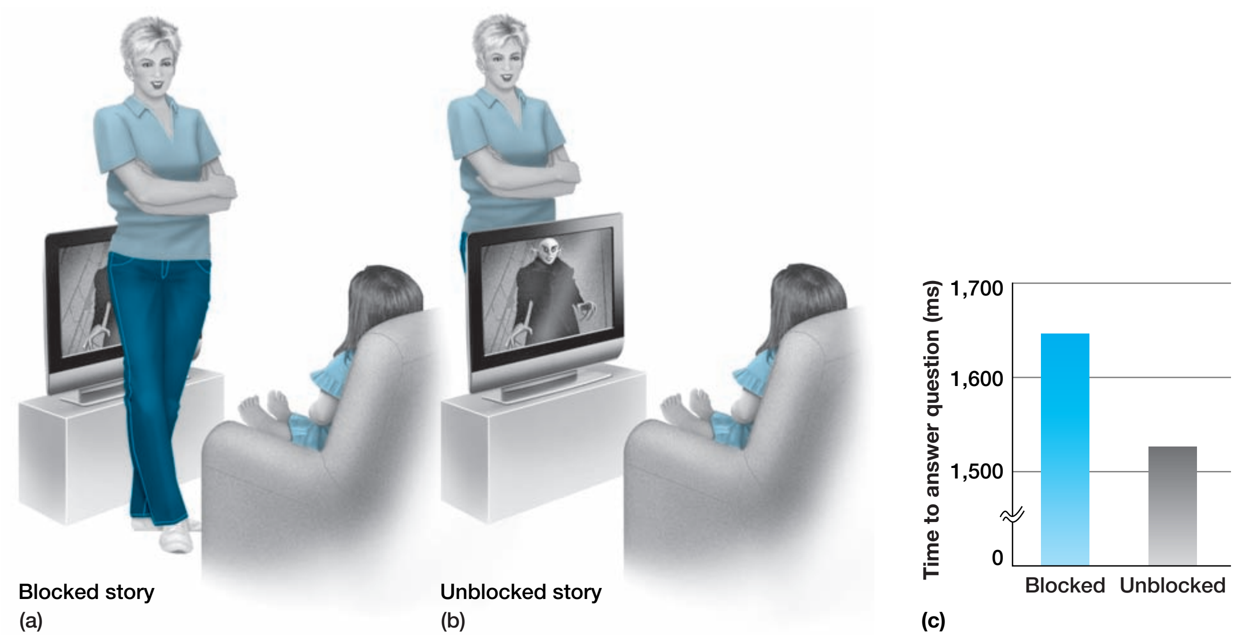

■ **Figure 10.11** Horton and Rapp's (2003) experiment. (a) In the blocked condition, the story describes Melanie's mother as being in front of the TV. (b) In the unblocked condition, the story describes Melanie's mother as being behind the TV. (c) The results indicate that the reaction time to answer a question about something happening on the screen is slower for the blocked condition.

contained the answer, was blocked, so Mary couldn't see it. The result, shown in Figure 10.11c, confirms this prediction—responding was slower in the blocked condition. This supports the idea that readers represent story events in a manner similar to actual perception. That is, they experience a story as if they are experiencing the situation described in the text.

Producing Language: Speech Errors

Up until now our emphasis has been on comprehending language—either reading text or listening to another person speak. But speech is usually not a one-way street; we are both listeners and speakers. This means we need to broaden our discussion to include not only how listeners understand the meaning of what others are saying, but also how speakers create meaning for others to understand.

Producing language is another one of those cognitive processes we achieve rapidly and easily, but which is actually extremely complex. This complexity becomes apparent when we consider that the act of speaking involves assembling strings of words that have been retrieved from memory. This retrieval is rapid—more than 3 words per second for normal conversation—and is drawn from a lexicon of more than 50,000 words. Not only are the correct words rapidly retrieved, but they are produced in the correct order and combined with other words to create a grammatically correct sentence (Dell, 1995; Levelt, 2001).

One technique for determining how we achieve this feat is to use a technique that has proven valuable in understanding other cognitive processes—analyzing the types of errors that people make. **Speech errors**, which are also called "slips of the tongue," were made famous by Sigmund Freud, who suggested that slips of the tongue reflected the speaker's unconscious motivations (Freud, 1901). According to this idea, introducing a guest at a party by saying "It gives me great pleasure to prevent . . . ," when the speaker means to say "It gives me great pleasure to present . . . ," might be revealing the speaker's distaste for the guest.

Although slips of this kind, which have been called **Freudian slips**, do occur, there is little to support Freud's idea that all slips are caused by unconscious motivation. For one thing, of all the slips that occur in everyday speech, only a small fraction can be linked to a person's unconscious motivations. In addition, this kind of explanation not only involves guessing as to what the unconscious motivation might be, but it is difficult or impossible to make predictions of how a particular person's unconscious motivations might become translated into speech errors. Thus, stories such as the one about introducing the guest, or the man who explained his excellent memory by saying, "I have a pornographic memory" can be "explained" in terms of unconscious motivation, but only after the fact (Dell, 1995).

Rather than focusing on unconscious motivation, speech researchers have used speech errors to provide insights into basic mechanisms of language (Bock, 1995; Dell, 1995; Garrett, 1975, 1980). However, studying speech errors is complicated by the fact that they are very rare, occurring only about 1 or 2 times out of every 1,000 words in normal conversation (Dell, 1995). Thus, one of the challenges of studying speech errors is to identify them. Researchers have used a number of methods to do this.

Method

Identifying Speech Errors

One way to identify speech errors is to note them as they happen in everyday speech. This requires great vigilance by the researcher, who must always be ready to write down an error when it occurs, and also a great deal of patience, because errors happen so infrequently. Another problem with this method is that it may result in a biased sample because errors that are funny or bizarre (for example, switching first letters to create "Hissing his mystery lecture") are more likely to be noticed. One way to avoid this sampling problem is to tape-record samples of speech and then carefully analyze the tapes for speech errors. When Garnham and coworkers (1981) did this, they identified 200 errors in 200,000 words.

Speech errors have also been created in the laboratory by rapidly presenting word pairs and using a tone to instruct participants to repeat the pair they just heard (Baars et al., 1975). This technique has the advantage of increasing the error rate to about 10 percent and also makes it possible to ask some specific questions about the conditions that result in speech errors. For example, using this technique, Baars and coworkers (1975) found that slips that create nonwords, like "beal dack" (when "deal back" was intended) are three times less likely than ones that create meaningful words, like "real dead" (when "deal red" was intended). A disadvantage of this laboratory technique is that it creates errors artificially. It is therefore important to collect both laboratory-produced errors and errors that occur naturally.

Speech Errors and Language Mechanisms

What do speech errors tell us about the basic mechanisms of language? To answer this question, language researchers focus on two aspects of speech errors:

1. *Frequency of different types of errors.* The most common errors indicate the basic units of language production. For example, two of the most common exchanges, which we will describe below, involve phonemes and words. The high frequency of these types of slips, plus other evidence, have led to the conclusion that phonemes and words are basic units of language (Dell, 1995).
2. *Patterns of errors.* Speech errors do not occur randomly. In many cases, researchers have identified rules that govern speech errors. These rules provide insights into mechanisms of normal language production. *Phoneme exchanges* and *word exchanges* illustrate two rules of speech errors.

Phoneme Exchanges **Phoneme exchanges**, such as saying "fleaky squoor" instead of "squeaky floor" (Fromkin, 1973), illustrate the **consonant-vowel rule**: Phonemes of the same type replace one another. Consonants replace consonants and vowels replace vowels. (This particular example involves the exchange of consonant clusters "sq" and "fl.")

Word Exchanges The following are two examples of **word exchanges**:

1. "I have to fill up my *gas* with *car*" instead of "I have to fill up my *car* with *gas*" (noun exchange).
2. "Once I *stop* I can't *start*" instead of "Once I *start* I can't *stop*" (verb exchange).

Word exchanges such as these follow the **syntactic category rule**: When one word replaces another, the same syntactic categories are used. Nouns slip with nouns (example 1), and verbs slip with verbs (example 2).

These examples indicate that speech errors are far from random. Speech errors follow rules that reflect the importance of specific sound units (consonants and vowels) and parts of speech (nouns and verbs). Our final example of speech errors, *word substitution*, illustrates how speech can be influenced by knowledge that a speaker brings to the situation.

Word Substitutions An example of **word substitution** is when someone says "Liszt's second Hungarian restaurant" instead of "Liszt's second Hungarian rhapsody." This type of error is probably caused by the fact that both "restaurant" and "rhapsody" are associated with Hungarian. This is therefore an example of an error based on the speaker's knowledge of classical music and ethnic restaurants. This is also an example of the syntactic category rule, because both restaurant and rhapsody are nouns. It also illustrates substitution of one three-syllable word beginning with "r" with another three-syllable word beginning with "r." Clearly, speech errors are influenced by numerous factors, related both to the basic structure of language and to a person's prior knowledge.

Producing Language: Conversations

Although language can be produced by a single person talking alone, as when a person recites a monologue or gives a speech, the most common form of language production is conversation—two or more people talking with one another. Conversation, or dialogue, provides another example of a cognitive skill that seems easy but contains underlying complexities.

In a conversation, other people are involved, so each person needs to take what the other people are saying into account (Pickering & Garrod, 2004). This is an impressive accomplishment because we often do not know what the other person is going to say. Nonetheless, we are usually able to respond to their statements almost immediately. One way that people deal with these difficulties is that they coordinate their conversations on both semantic and syntactic levels.

Semantic Coordination

When people are talking about a topic, each person brings his or her own knowledge to the conversation, and conversations go more smoothly when the participants bring *shared* knowledge to a conversation. Thus, when people are talking about current events, it helps when everyone has been keeping up with the news, and is more difficult when one of the people has just returned from 6 months of meditation in an isolated monastery.

But even when everyone brings similar knowledge to a conversation, it helps when speakers take steps to guide their listeners through the conversation. One way this can be achieved is by following the *given–new contract*. The **given–new contract** states that the speaker should construct sentences so they include two kinds of information: (1) *given information*—information that the listener already knows; and (2) *new information*—information that the listener is hearing for the first time (Haviland & Clark, 1974). For example, consider the following two sentences.

1. Ed was given an alligator for his birthday.
   Given information (from previous conversation): Ed had a birthday.
   New information: He got an alligator.
2. The alligator was his favorite present.
   Given information (from sentence 1): Ed got an alligator.
   New information: It was his favorite present.

Notice how the new information in the first sentence becomes the given information in the second sentence. Susan Haviland and Herbert Clark (1974) demonstrated the consequences of not following the given–new contract by presenting pairs of sentences and asking participants to press a button when they felt they understood the second sentence in each pair. They found that it took longer for the participants to comprehend the second sentence in pairs like this one:

We checked the picnic supplies.
The beer was warm.

than it took to comprehend the second sentence in pairs like this one:

We got some beer out of the trunk.
The beer was warm.

The reason comprehending the second sentence in the first pair takes longer is that the given information, *that there were picnic supplies*, does not mention beer. Thus, the reader or listener needs to make an inference in the first case that beer was among the picnic supplies. In contrast, this inference is not required in the second pair because the first sentence includes the information that beer is in the trunk.

Syntactic Coordination

When two people exchange statements in a conversation, it is common for them to use similar grammatical constructions. Kathryn Bock (1990) provides the following example, taken from a recorded conversation between a bank robber and his lookout, which was intercepted by a ham radio operator as the robber was removing the equivalent of a million dollars from a bank vault in England.

Robber: "*. . . you've got to hear* and witness it *to realize how bad it is*."
Lookout: "You *have got to experience exactly* the same position as me, mate, *to understand how I feel*." (from Schenkein, 1980, p. 22)

Bock has added italics to the statements to illustrate how the lookout has copied the form of the robber's statement. This copying of form reflects a phenomenon called **syntactic priming**—hearing a statement with a particular syntactic construction increases the chances that a sentence will be produced with the same construction. Syntactic priming is important because it can lead people to coordinate the grammatical form of their statements during a conversation. Holly Branigan and coworkers (2000) illustrated syntactic priming by using the following procedure to set up a give-and-take between two people.

Method

Syntactic Priming

In a syntactic priming experiment two people engage in a conversation, and the experimenter determines whether production of a specific grammatical construction by one person increases the chances that the same construction will be used by the other person.

Participants in Branigan's experiment were told that the experiment was about how people communicate when they can't see each other. They thought they were working with another participant who was on the other side of a screen (the left person in Figure 10.12a). In reality, the person on the other side of the screen was a confederate who was working with the experimenter.

The confederate began the experiment by making a *priming statement,* as shown on the left of Figure 10.12a. This statement was in one of the following two forms:

1. "The girl gave the book to the boy."
2. "The girl gave the boy the book."

The participant had two tasks: (1) find the *matching card* on the table that corresponded to the confederate's statement, as shown on the right of Figure 10.12a; and (2) describe one of the *response cards* to the confederate, as shown in Figure 10.12b. We can conclude that syntactic priming has occurred if the form of the participant's statement to the confederate matches the form of the confederate's original statement.

Branigan found that on 78 percent of the trials, the form of the participant's statement matched the form of the confederate's priming statement. Thus, if the participant heard the confederate say "The girl gave the boy the book," this increased the chances that the participant would describe a response card like the one shown in Figure 10.12b as "The father brought his daughter a present" (instead of "The father brought a present for his daughter" or some other construction). This supports the idea that speakers are sensitive to the linguistic behavior of other speakers and adjust their behaviors to match. This coordination of syntactic form between speakers reduces the computational load involved in creating a conversation because it is easier to copy the form of someone else's sentence than it is to create your own form from scratch.

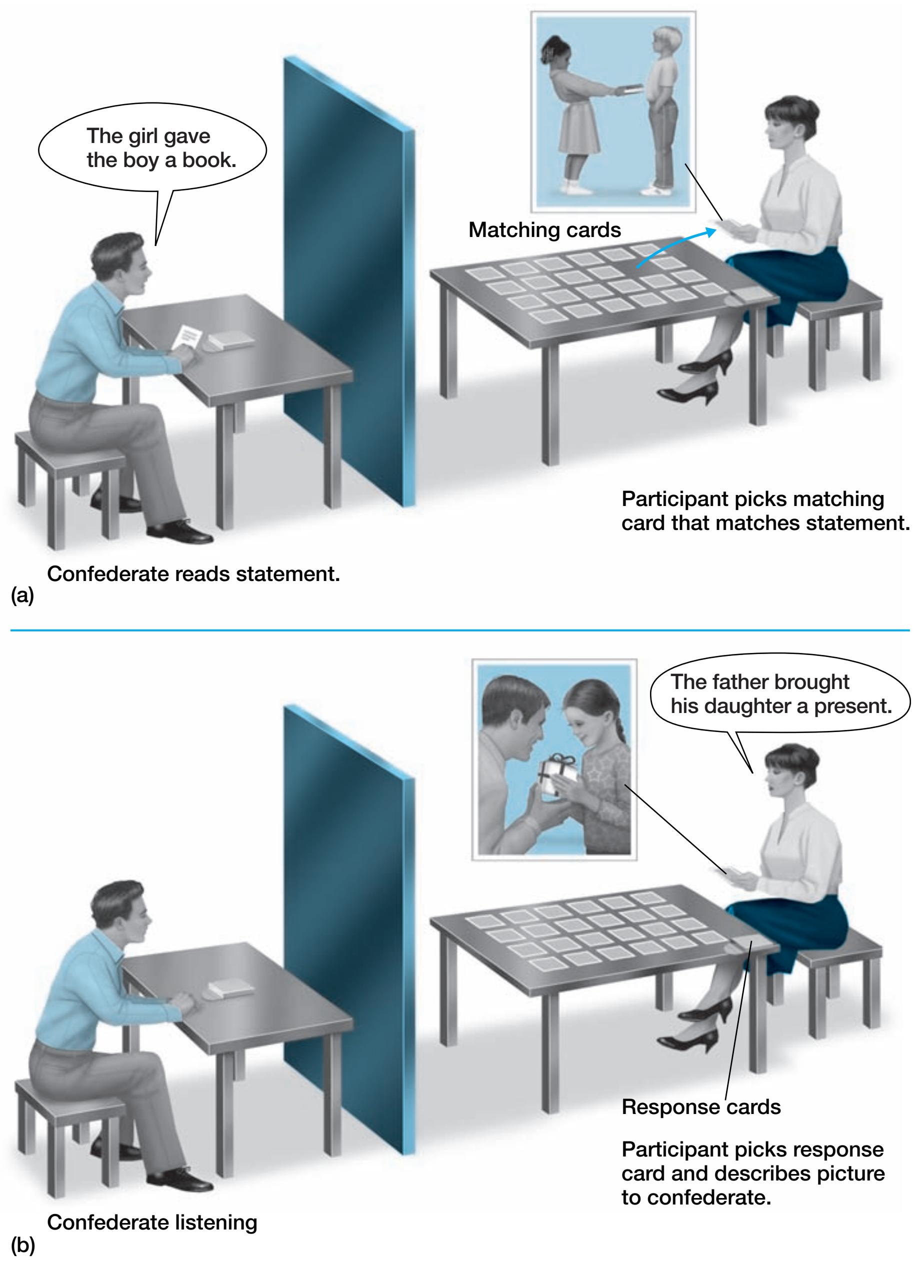

■ **Figure 10.12** The Branigan et al. (2000) experiment. (a) The participant, on the right, picks from the matching cards on the table a card with a picture that matches the statement read by the confederate. (b) The participant then picks a card from the pile of response cards and describes the picture on the response card to the confederate.

Let's summarize what we have said about conversations: Conversations are dynamic and rapid, but a number of processes make them easier. On the semantic side, people take other people's knowledge into account (if they don't, confusion can result). On the syntactic side, people coordinate or align the syntactic form of their statements. This makes speaking easier and frees up resources to deal with the tasks of alternating between understanding and producing messages that is the hallmark of successful conversations.

Something to Consider

Culture, Language, and Cognition

According to the **Sapir-Whorf hypothesis**, which was proposed by anthropologist Edward Sapir and linguist Benjamin Whorf, the nature of a culture's language can affect the way people think (Whorf, 1956). Although there was little evidence to support this when Whorf made his proposal, recent research has provided evidence that favors the idea that language can influence cognition. Debi Roberson and coworkers (2000; also see Davidoff, 2001) demonstrated that language can affect the way people perceive colors, by comparing color perception in British participants and participants from a culture called the Berinmo from New Guinea.

Roberson had the British and Berinmo participants name the colors of 160 Munsell color chips (small color chips similar to those found in paint stores, but scientifically color-calibrated to be used in research). The results for the British and Berinmo are shown in ● Color Plate 10.1. These diagrams indicate the color names that each group assigned to each of the chips. For example, chip 5B-9, indicated by the dot in the diagrams, was called *blue* by the British participants and *wap* by the Berinmo participants.

One difference in the results for the two cultures was that the British used eight different color names (blue, green, yellow, pink, red, brown, orange, and purple) and the Berinmo used just five (wap, wor, mehi, kel, and nol). Another difference was that the borders between the colors were different. For example, look at the area marked *yellow* for the British, and compare this to *wor* for the Berinmo. The Berinmo gave one name (wor) to many of the chips that the British called yellow, orange, green, and brown.

Having demonstrated differences in how participants in the two cultures *assigned names* to the color chips, Roberson and coworkers did an additional experiment to determine whether the two *perceive* colors differently. They accomplished this by doing a categorical perception experiment.

Method

Categorical Perception

In a categorical perception experiment, participants are presented with pairs of stimuli and are asked to indicate whether they are the same or different. For example, would you say chips A and B in ● Color Plate 10.2 are the same or different? How about chips B and C?

CogLab
Categorical Perception—Identification

It was probably easier to make the judgment "different" for chips B and C. This is because most English speakers place chips A and B in the same category (green), but place B and C in different categories (green and blue). **Categorical perception** occurs when stimuli that are in the same categories (like A and B) are more difficult to discriminate from one another than are stimuli that are in different categories (like B and C).

CogLab
Categorical Perception—Discrimination

When Roberson and coworkers presented the British and Berinmo participants with pairs of colors and asked them to indicate whether they were the same or different, British participants discriminated more easily between blue and green chips than the Berinmo, but the Berinmo discriminated more easily between *nol* and *wor* than the British. This means that differences in the way names were assigned to colors (indicated by the diagrams in ● Color Plate 10.1) affected the ability to tell the difference between colors. The fact that language has an effect on behavior supports the Sapir-Whorf hypothesis (see also Gentner & Goldwin-Meadow, 2003).

Although Roberson's experiments show that language can affect color perception, there is also evidence for similarities in color perception across different languages. Terry Regier and coworkers (2005) provided this evidence by tabulating the color naming data from speakers of over 100 different languages. Each participant was presented with the chips in ● Color Plate 10.3 and was asked to name the color of each chip. This is what Roberson did, but in addition participants were asked to pick the *best* example of each color. For example, English speakers would pick the best red, the best green, and so on. Regier's results, shown in ● Color Plate 10.1c, indicate that the "best" colors tend to cluster around the areas that English speakers call red, yellow, green, and blue. Notice that there is some variation, but that there are four distinct "islands" that correspond to the best examples of each color.

What do these results mean? Apparently, language can affect color perception, as Roberson showed, but there are also limits to the effects of language, as Regier showed. Other experiments have demonstrated differences in how Westerners and East Asians think about objects (Iwao & Gentner, 1997), numbers (Lucy & Gaskins, 1997), and space (Levinson, 1996). See If You Want to Know More on page 392 for more references on the connection between language and thinking.

Test Yourself 10.2

1. Why do we say that understanding a story involves more than adding up the meanings of the sentences that make up the story?
2. What is coherence? Inference? What are the different types of inference, and what is their relation to coherence?
3. What are the assumptions behind the situation model, and what predictions does this model make?

4. What are speech errors? Describe Freud's ideas about speech errors and why modern language researchers do not consider these ideas to be that important.
5. What aspects of speech errors provide information about language mechanisms? What do phoneme exchanges and word exchanges tell us about language? What does the example of word substitution that involves Hungarian restaurants indicate about language mechanisms?
6. Describe how semantic coordination and syntactic coordination facilitate conversations. Be sure you understand syntactic priming and what it demonstrates about language production.
7. What is the Sapir-Whorf hypothesis? Describe experiments on color perception that support this hypothesis. Also describe the evidence that indicates that some aspects of color perception are the same across different languages.

Chapter Summary

1. Language is a system of communication that uses sounds or symbols that enables us to express our feelings, thoughts, ideas, and experiences. Human language can be distinguished from animal communication by its creativity, hierarchical structure, governing rules, and universality.
2. Modern research in the psychology of language blossomed in the 1950s and 1960s, with the advent of the cognitive revolution. One of the central events in the cognitive revolution was Chomsky's critique of Skinner's behavioristic analysis of language.
3. All the words a person knows are his or her lexicon. Phonemes and morphemes are two basic units of words.
4. The effect of meaning on the perception of phonemes is illustrated by the phonemic restoration effect. Meaning, as well as a person's experience with other aspects of language, are important for achieving speech segmentation.
5. The ability to understand words is influenced by word frequency and the context provided by the sentence.
6. Lexical ambiguity refers to the fact that a word can have more than one meaning and that the word's meaning in a sentence may not be clear. Lexical priming experiments show that all meanings of a word are activated immediately after the word is presented, but then context determines the eventual meaning of the word.
7. The meaning of a sentence is determined by both semantics (the meanings of words) and syntax (rules for using words in sentences).
8. Parsing is the process by which words in a sentence are grouped into phrases. Grouping into phrases is a major determinant of the meaning of a sentence. This process has been studied by using ambiguous sentences.
9. Two mechanisms proposed to explain parsing are (1) the syntax-first approach and (2) the interactionist approach. The syntax-first approach emphasizes how syntactic principles such as late closure determine how a sentence is parsed. The interaction-

ist approach states that both semantics and syntax operate simultaneously to determine parsing. This approach is supported by eye-movement studies.

10. Coherence enables us to understand stories. Coherence is largely determined by inference. Three major types of inference are anaphoric, instrumental, and causal.
11. The situation model approach to text comprehension states that people represent the situation in a story in terms of the people, objects, locations, and events that are being described in the story. Experiments support the idea that a reader often takes the protagonist's point of view in the story.
12. One of the major tools in the study of speech production is the determination and interpretation of speech errors (slips of the tongue). Research showing that some speech errors are more common than others and that speech errors often follow rules have provided insights into the mechanisms of normal language production.
13. Conversations, which involve give and take between two or more people, are made easier by two mechanisms of cooperation between participants in a conversation—semantic coordination and syntactic coordination. Syntactic priming experiments provide evidence for syntactic coordination.
14. There is evidence that a culture's language can influence the way people perceive and think. Experiments comparing color perception in Westerners and people in the Berinmo culture have revealed differences in color perception related to language. However, there is also evidence for some consistency in color perception across different languages.

Think About It

1. How do the ideas of coherence and connection apply to some of the movies you have seen lately? Have you found that some movies are easy to understand whereas others are more difficult? In the movies that are easy to understand, does one thing appear to follow from another, whereas in the more difficult ones, some things seem to be left out? What is the difference in the "mental work" needed to determine what is going on in these two kinds of movies? (You can also apply this kind of analysis to books you have read.)
2. Next time you are able to eavesdrop on a conversation, notice how the give-and-take among participants follows (or does not follow) the given–new contract. Also, notice how people change topics and how that affects the flow of the conversation. Finally, see if you can find any evidence of syntactic priming. One way to "eavesdrop" is to be part of a conversation that includes at least two other people. But don't forget to say something every so often!
3. One of the interesting things about languages is the use of "figures of speech," which people who know the language understand but which nonnative speakers often find baffling. One example is the sentence "He brought everything but the kitchen sink." Can you think of other examples? If you speak a language other than

English, can you identify figures of speech in that language that might be baffling to English speakers?

4. Newspaper headlines are often good sources of ambiguous phrases. For example, consider the following, which were actual headlines: “Milk drinkers are turning to powder,” “Iraqi head seeks arms,” “Farm bill dies in house,” and “Squad helps dog bite victim.” See if you can find examples of ambiguous headlines in the newspaper, and try to figure out what it is that makes the headlines ambiguous.
5. People often say things in an indirect way, but listeners can often still understand what they mean. See if you can detect these indirect statements in normal conversation. (Examples: “Do you want to turn left here?” to mean “I think you should turn left here”; “Is it cold in here?” to mean “Please close the window.”)

If You Want to Know More

1. Animal language. Can monkeys use language in a way similar to humans? This is a controversial question, with some psychologists answering “yes” and others “no.”

Savage-Rumbaugh, S., & Lewin, R. (1994). *Kanzi, the ape at the brink of the human mind.* New York: Wiley.

2. Indirect statements. People use indirect statements all the time (see preceding Think About It). There is evidence that indirect statements are more prevalent in some cultures than in others.

Holtgraves, T. (1998). Interpreting indirect replies. *Cognitive Psychology*, *37*, 1–27.

3. Bilingualism. When people speak two or more languages, are these languages stored together or separately? This question, as well as other questions about the mechanisms involved in bilingualism, has been studied both behaviorally and physiologically.

Kroll, J. F., & Tokowicz, N. (2005). Models of bilingual representation and processing: Looking back and to the future. In J. F. Kroll & A. M. B. De Groot (Eds.), *Handbook of bilingualism: Psycholinguistic approaches* (pp. 531–553). New York: Oxford University Press.

Perani, D., & Abutalebi, J. (2005). The neural basis of first and second language processing. *Current Opinion in Neurobiology*, *15*, 202–206.

Petitto, L. A., Katerelos, M., Levy, B. G., Gauna, K., Tétreault, K., & Ferraro, V. (2001). Bilingual signed and spoken language acquisition from birth: Implications for the mechanisms underlying early bilingual language acquisition. *Journal of Child Language*, *28*, 453–496.

Snow, C. E. (1998). Bilingualism and second language acquisition. In J. B. Gleason & N. B. Ratner (Eds.) *Psycholinguistics* (2nd ed., pp. 453–481). Ft. Worth, TX: Harcourt.

4. Psychology of reading. Much of our use of language involves reading. This involves vision or touch (in the case of Braille) and demands on memory that are different than for spoken language.

Price, C. J., & Mechelli, A. (2005). Reading and reading disturbance. *Current Opinion in Neurobiology*, *15*, 231–238.

Starr, M. S., & Rayner, K. (2001). Eye movements during reading: Some current controversies. *Trends in Cognitive Sciences*, *5*, 156–163.

5. Language, culture, and the representation of space. Figure 10.13 indicates three ways of expressing spatial relationships. Different cultures favor different systems, and there is evidence that language plays an important role in this.

Majid, M., Bowerman, M., Kita, S., Haun, D. B. M., & Levinson, S. C. (2004). Can language restructure cognition? *Trends in Cognitive Sciences*, *8*, 108–114.

6. Culture and categories. Which two objects in Figure 10.14 would you place together? Which two of the following words would you place together? Panda, Monkey, Banana. Research has shown that Chinese and Americans sort these items differently and that these differences may be related to language.

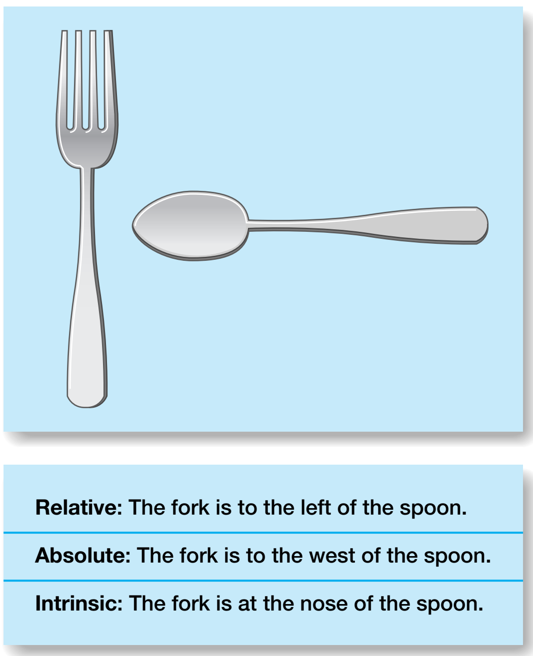

■ **Figure 10.13** Three ways of expressing spatial relationships (Majid et al., 2004).

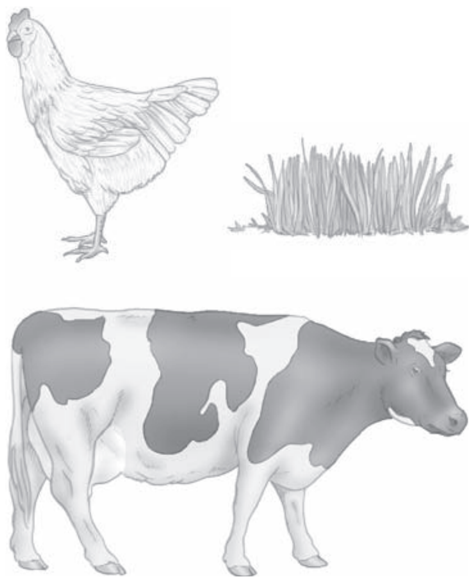

■ **Figure 10.14** Which objects belong together? (Based on Chiu, 1972.)

Chiu, L-H. (1972). A cross-cultural comparison of cognitive styles in Chinese and American children. *International Journal of Psychology*, 7, 235–242.

Ji, L., Peng, K., & Nisbett, R. E. (2000). Culture, control, and perception of relationships in the environment. *Journal of Personality and Social Psychology*, *78*, 943–955.

7. Effect of brain damage on language. In the 1800s Paul Broca and Karl Wernicke identified areas in the frontal and temporal lobes of the brain that when damaged cause aphasia—disorders of language. Modern researchers have identified many types of aphasia.

Farah, M. J., & Feinberg, T. E. (2000). *Patient-based approaches to cognitive neuroscience* (pp. 165–271). Cambridge, MA: MIT Press.

8. Language acquisition. Children usually begin learning language before they can speak, produce their first words at about a year, and have mastered the basic linguistic structures of language by about 3 years of age.

Carroll, D. W. (2004). *Psychology of language* (4th ed.). Belmont, CA: Wadsworth.

Gleason, J. B., & Ratner, N. B. (1998). Language acquisition. In J. B. Gleason & N. B. Ratner (Eds.), *Psycholinguistics* (2nd ed., pp. 347–407). Fort Worth, TX: Harcourt.

Key Terms

Anaphoric inference, 377
Categorical perception, 388
Causal inference, 378
Coherence, 377
Consonant-vowel rule, 382
Freudian slip, 381
Garden-path model, 372
Garden-path sentence, 372
Given–new contract, 384
Inference, 377
Instrumental inference, 378
Interactionist approach to parsing, 373
Language, 357
Late closure, 371
Lexical ambiguity, 365
Lexical decision task, 364
Lexical priming, 366
Lexicon, 360
Morpheme, 361
Parser, 370
Parsing, 369
Phoneme, 361
Phoneme exchange, 382
Phonemic restoration effect, 361
Psycholinguistics, 359
Sapir-Whorf hypothesis, 387
Semantics, 368
Situation model, 379
Speech error, 381
Speech segmentation, 363
Syntactic ambiguity, 370
Syntactic category rule, 383
Syntactic priming, 385
Syntax, 368
Syntax-first approach to parsing, 371
Word exchange, 383
Word frequency, 363
Word-frequency effect, 364
Word substitution, 383

CogLab

To experience these experiments for yourself, go to http://coglab.wadsworth.com. Be sure to read each experiment's setup instructions before you go to the experiment itself. Otherwise, you won't know which keys to press.

Primary Labs

Word superiority How speed of identifying a letter compares when the letter is isolated or in a word (p. 361).

Lexical decision Demonstration of the lexical decision task, which has been used to provide evidence for the concept of spreading activation (p. 364).

Categorical perception—identification Demonstration of categorical perception based on the identification of different sound categories (p. 388).

Categorical perception—discrimination Demonstration of categorical perception based on the ability to discriminate between sounds (p. 388).

11 Problem Solving

What Is a Problem?

The Gestalt Approach: Problem Solving as Representation and Restructuring

- Representing a Problem in the Mind
- Insight in Problem Solving
- Demonstration: Two Insight Problems
- Obstacles to Problem Solving
- Demonstration: The Candle Problem

Modern Research on Problem Solving: The Information-Processing Approach

- Newell and Simon's Approach
- Demonstration: Tower of Hanoi Problem
- The Importance of How a Problem Is Stated
- Demonstration: The Mutilated-Checkerboard Problem
- Method: Think-Aloud Protocol

Test Yourself 11.1

Using Analogies to Solve Problems

- Method: Analogical Transfer
- Analogical Problem Solving and the Duncker Radiation Problem
- Demonstration: Duncker's Radiation Problem
- Analogical Encoding
- Analogy in the Real World
- Method: In-vivo Problem-Solving Research

How Experts Solve Problems

- Some Differences Between How Experts and Novices Solve Problems

Creative Problem Solving

- Demonstration: Creating an Object

Something to Consider: Sleep Inspires Insight

Test Yourself 11.2

Chapter Summary

Think About It

If You Want to Know More

Key Terms

Some Questions We Will Consider

- How does the ability to solve a problem depend on how the problem is represented in the mind? (397)
- Is there anything special about "insight" problems? (398)
- How can analogies be used to help solve problems? (413)
- What is the difference between how experts in a field solve problems and how nonexperts solve problems? (421)