SOME QUESTIONS WE WILL CONSIDER

- Is it possible to focus attention on just one thing, even when there are many other things going on at the same time? (95)
- Under what conditions can we pay attention to more than one thing at a time? (110)
- What does attention research tell us about the effect of talking on cell phones while driving a car? (112)
- Is it true that we are not paying attention to a large fraction of the things happening in our environment? (115)

Roger, sitting in the library, is attempting to do his math homework when some people at the next table start talking. He is annoyed because people aren't supposed to talk in the library, but he is so focused on the math problems that it doesn't distract him (**Figure 4.1a**). However, a little later, when he decides to take a break from his math homework and play an easy game on his cell phone, he does find their conversation distracting (**Figure 4.1b**). "Interesting," he thinks. "Their talking didn't bother me when I was doing the math problems."

Deciding to stop resisting his listening to the conversation, Roger begins to consciously eavesdrop while continuing to play his cell phone game (**Figure 4.1c**). But just as he is beginning to figure out what the couple is talking about, his attention is captured by a loud noise and commotion from across the room, where it appears a book cart has overturned, scattering books on the floor. As he notices that one person seems upset and others are gathering up the books, he looks from one person to another and decides he doesn't know any of them (**Figure 4.1d**).

Figure 4.1 Roger's adventures with attention. (a) Selective attention: doing math problems while not being distracted by people talking. (b) Distraction: playing a game but being distracted by the people talking. (c) Divided attention: playing the game while listening in on the conversation. (d) Attentional capture and scanning: a noise attracts his attention, and he scans the scene to figure out what is happening.

Roger's experiences illustrate different aspects of **attention**—the ability to focus on specific stimuli or locations. His attempt to focus on his math homework while ignoring the people talking is an example of **selective attention**—attending to one thing while ignoring others. The way the conversation in the library interfered with his cell phone game is an example of **distraction**—one stimulus interfering with the processing of another stimulus. When Roger decides to listen in on the conversation while simultaneously playing the game, he is displaying **divided attention**—paying attention to more than one thing at a time. Later, his eavesdropping is interrupted by the noise of the overturned book cart, an example of **attentional capture**—a rapid shifting of attention usually caused by a stimulus such as a loud noise, bright light, or sudden movement. Finally, Roger's attempt to identify the people across the room, looking from one person's face to another, is an example of **visual scanning**—movements of the eyes from one location or object to another.

With all of these different aspects of attention in mind, let's return to William James's (1890) definition of attention, which we introduced in Chapter 1:

> Millions of items … are present to my senses which never properly enter my experience. Why? Because they have no interest for me. My experience is what I agree to attend to. . . . Everyone knows what attention is. It is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. . . . It implies withdrawal from some things in order to deal effectively with others.

Although this definition is considered a classic, and certainly does capture a central characteristic of attention—withdrawal from some things in order to deal effectively with others—we can now see that it doesn't capture the diversity of phenomena that are associated with attention. Attention, as it turns out, is not one thing. There are many different aspects of attention, which have been studied using different approaches.

This chapter, therefore, consists of a number of sections, each of which is about a different aspect of attention. We begin with a little history, because early research on attention helped establish the information processing approach to cognition, which became the central focus of the new field of cognitive psychology.

Attention as Information Processing

Modern research on attention began in the 1950s with the introduction of Broadbent's filter model of attention.

Broadbent's Filter Model of Attention

Broadbent's **filter model of attention**, which we introduced in Chapter 1 (page 14), was designed to explain the results of an experiment done by Colin Cherry (1953). Cherry studied attention using a technique called **dichotic listening**, where *dichotic* refers to presenting different stimuli to the left and right ears. The participant's task in this experiment is to focus on the message in one ear, called the attended ear, and to repeat what he or she is hearing out loud. This procedure of repeating the words as they are heard is called **shadowing** (**Figure 4.2**).

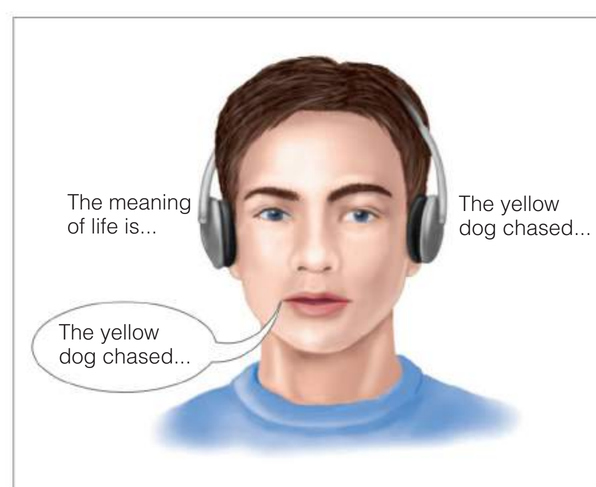

Figure 4.2 In the shadowing procedure, which involves dichotic listening, a person repeats out loud the words that they have just heard. This ensures that participants are focusing their attention on the attended message.

Cherry found that although his participants could easily shadow a spoken message presented to the attended ear, and they could report whether the unattended message was spoken by a male or female, they couldn't report what was being said in the unattended ear. Other dichotic listening experiments confirmed that people are not aware of most of the information being presented to the unattended ear. For example, Neville Moray (1959) showed that participants were unaware of a word that had

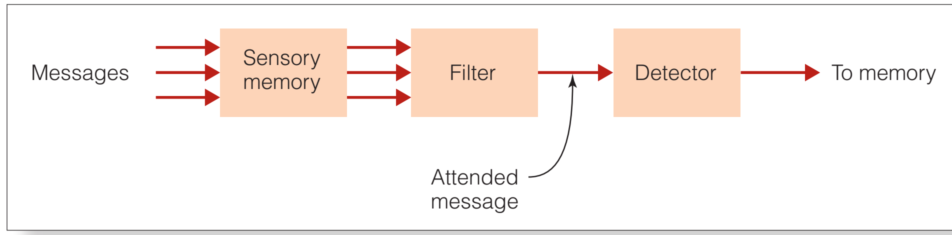

➤ **Figure 4.3** Flow diagram of Broadbent's filter model of attention.

been repeated 35 times in the unattended ear. The ability to focus on one stimulus while filtering out other stimuli has been called the **cocktail party effect**, because at noisy parties people are able to focus on what one person is saying even if there are many conversations happening at the same time.

Based on results such as these, Donald Broadbent (1958) created a model of attention designed to explain how it is possible to focus on one message and why information isn't taken in from the other message. This model, which introduced the flow diagram to cognitive psychology (see page 14), proposed that information passes through the following stages (**Figure 4.3**):

- *Sensory memory* holds all of the incoming information for a fraction of a second and then transfers all of it to the filter. (We will discuss sensory memory in Chapter 5.)
- The **filter** identifies the message that is being attended to based on its physical characteristics—things like the speaker's tone of voice, pitch, speed of talking, and accent—and lets only this attended message pass through to the detector in the next stage. All of the other messages are filtered out.
- The **detector** processes the information from the attended message to determine higher-level characteristics of the message, such as its meaning. Because only the important, attended information has been let through the filter, the detector processes all of the information that enters it.
- The output of the detector is sent to *short-term memory*, which holds information for 10–15 seconds and also transfers information into *long-term memory*, which can hold information indefinitely. We will describe short- and long-term memory in Chapters 5–8.

Broadbent's model is called an **early selection model** because the filter eliminates the unattended information right at the beginning of the flow of information.

Modifying Broadbent's Model: More Early Selection Models

The beauty of Broadbent's filter model of attention was that it provided testable predictions about selective attention, which stimulated further research. One prediction is that since all of the unattended messages are filtered out, we should not be conscious of information in the unattended messages. To test this idea, Neville Moray (1959) did a dichotic listening experiment in which his participants were instructed to shadow the message presented to one ear and to ignore the message presented to the other ear. But when Moray presented the listener's name to the unattended ear, about a third of the participants detected it (also see Wood & Cowan, 1995).

Moray's participants had recognized their names even though, according to Broadbent's theory, the filter is supposed to let through only one message, based on its physical characteristics. Clearly, the person's name had not been filtered out, and, most important, it had been analyzed enough to determine its meaning. You may have had an experience similar to Moray's laboratory demonstration if, as you were talking to someone in a noisy room, you suddenly heard someone else say your name.

Following Moray's lead, other experimenters showed that information presented to the unattended ear is processed enough to provide the listener with some awareness of its meaning. For example, J. A. Gray and A. I. Wedderburn (1960), while undergraduates at the University of Oxford, did the following experiment, sometimes called the "Dear Aunt Jane" experiment. As in Cherry's dichotic listening experiment, the participants were told to shadow the message presented to one ear. As you can see in **Figure 4.4**, the attended (shadowed) ear received the message "Dear 7 Jane," and the unattended ear received the message "9 Aunt 6." However, rather than reporting the "Dear 7 Jane" message that was presented to the attended ear, participants reported hearing "Dear Aunt Jane."

Switching to the unattended channel to say "Aunt" means that the participant's attention had jumped from one ear to the other and then back again. This occurred because they were taking the meaning of the words into account. (An example of top-down processing! See page 67.)

Because of results such as these, Anne Treisman (1964) proposed a modification of Broadbent's model. Treisman proposed that selection occurs in two stages, and she replaced Broadbent's filter with an attenuator (**Figure 4.5**). The **attenuator** analyzes the incoming message in terms of (1) its physical characteristics—whether it is high-pitched or low-pitched, fast or slow; (2) its language—how the message groups into syllables or words; and (3) its meaning—how sequences of words create meaningful phrases. Note that the attenuator represents a process and is not identified with a specific brain structure.

Treisman's idea that the information in the channel is selected is similar to what Broadbent proposed, but in Treisman's **attenuation model of attention**, language and meaning can also be used to separate the messages. However, Treisman proposed that the analysis of the message proceeds only as far as is necessary to identify the attended message. For example, if there are two messages, one in a male voice and one in a female voice, then analysis at the physical level (which Broadbent emphasized) is adequate to separate the low-pitched male voice from the higher-pitched female voice. If, however, the voices are similar, then it might be necessary to use meaning to separate the two messages.

According to Treisman's model, once the attended and unattended messages have been identified, both messages pass through the attenuator, but the attended message emerges at full strength and the unattended messages are attenuated—they are still present but are weaker than the attended message. Because at least some of the unattended message gets through the attenuator, Treisman's model has been called a "leaky filter" model.

The final output of the system is determined in the second stage, when the message is analyzed by the dictionary unit. The **dictionary unit** contains words, stored in memory, each of which has a threshold for being activated (**Figure 4.6**). A threshold is the smallest signal strength that can barely be detected. Thus, a word with a low threshold might be detected even when it is presented softly or is obscured by other words.

According to Treisman, words that are common or especially important, such as the listener's name, have low thresholds, so even a weak signal in the unattended channel can activate that word, and

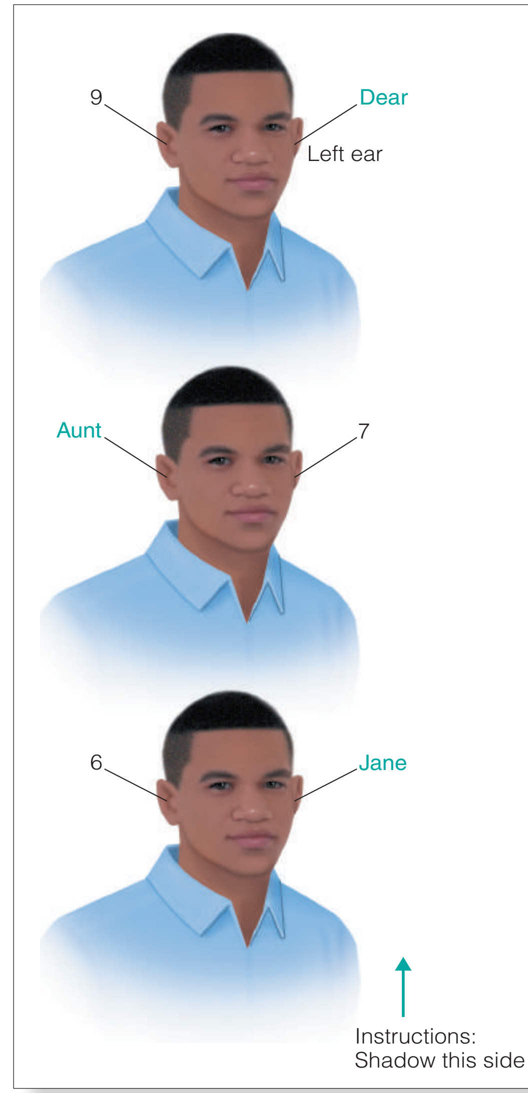

➤ **Figure 4.4** In Gray and Wedderburn's (1960) "Dear Aunt Jane" experiment, participants were told to shadow the message presented to the left ear. But they reported hearing the message "Dear Aunt Jane," which starts in the left ear, jumps to the right ear, and then goes back to the left ear.

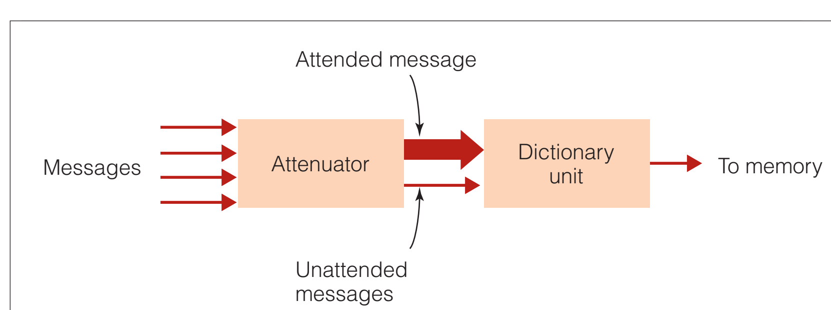

➤ **Figure 4.5** Flow diagram for Treisman's attenuation model of selective attention.

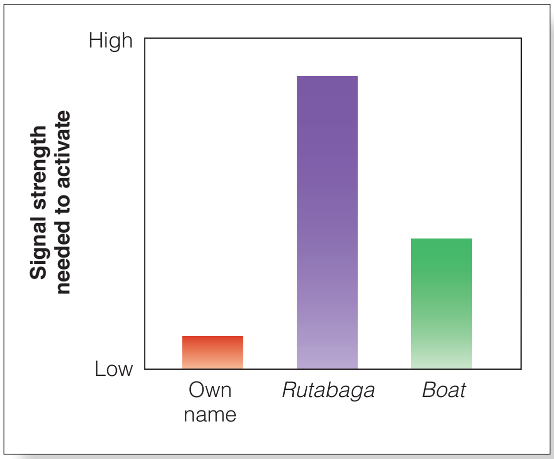

➤ **Figure 4.6** The dictionary unit of Treisman's attenuation model of selective attention contains words, each of which has a threshold for being detected. This graph shows the thresholds that might exist for three words. The person's name has a low threshold, so it will be easily detected. The thresholds for the words *rutabaga* and *boat* are higher, because they are used less or are less important to this particular listener.

we hear our name from across the room. Uncommon words or words that are unimportant to the listener have higher thresholds, so it takes the strong signal of the attended message to activate these words. Thus, according to Treisman, the attended message gets through, plus some parts of the weaker, unattended messages.

The research we have been describing so far was extremely important, not only because it defined some of the basic phenomena of attention but also because it demonstrated how an aspect of cognition could be conceptualized as a problem of information processing, in which information from the environment passes through various stages of processing. Like Broadbent's model, Treisman's is called an early selection model because it proposes a filter that operates at an early stage in the flow of information. Other models propose that selection can occur later.

A Late Selection Model

Other theories were proposed to take into account the results of experiments showing that messages can be selected at a later stage of processing, based primarily on their meaning. For example, in an experiment by Donald MacKay (1973), a participant listened to an ambiguous sentence, such as "They were throwing stones at the bank," that could be interpreted in more than one way. (In this example, "bank" can refer to a riverbank or to a financial institution.) These ambiguous sentences were presented to the attended ear while biasing words were presented to the other, unattended ear. For example, as the participant was shadowing "They were throwing stones at the bank," either the word "river" or the word "money" was presented to the unattended ear.

After hearing a number of ambiguous sentences, the participants were presented with pairs of sentences, such as "They threw stones toward the side of the river yesterday" and "They threw stones at the savings and loan association yesterday," and asked to indicate which of these two sentences was closest in meaning to one of the sentences they had heard previously. MacKay found that the meaning of the biasing word affected the participants' choice. For example, if the biasing word was "money," participants were more likely to pick the second sentence. This occurred even though participants reported that they were unaware of the biasing words that had been presented to the unattended ear.

MacKay proposed that because the meaning of the word *river* or *money* was affecting the participants' judgments, the word must have been processed to the level of meaning even though it was unattended. Results such as this led MacKay and other theorists to develop **late selection models of attention**, which proposed that most of the incoming information is processed to the level of meaning before the message to be further processed is selected (Deutsch & Deutsch, 1963; Norman, 1968).

The attention research we have been describing has focused on when selective attention occurs (early or late) and what types of information are used for the selection (physical characteristics or meaning). But as research in selective attention progressed, researchers realized that there is no one answer to what has been called the "early–late" controversy. Early selection can be demonstrated under some conditions and later selection under others, depending on the observer's task and the type of stimuli presented. Thus, researchers began focusing instead on understanding the many different factors that control attention.

This brings us back to Roger's experience in the library. Remember that he was able to ignore the people talking when he was doing his math homework but became distracted by the talking when he was playing the easy cell phone game. The idea that the ability to selectively attend to a task can depend both on the distracting stimulus and on the nature of the task has been studied by Nilli Lavie (2010), who introduced the concepts of *processing capacity* and *perceptual load*.

Processing Capacity and Perceptual Load

How do people ignore distracting stimuli when they are trying to focus their attention on a task? Lavie answers this question by considering two factors: (1) **processing capacity**, which refers to the amount of information people can handle and sets a limit on their ability to process incoming information; and (2) **perceptual load**, which is related to the difficulty of a task. Some tasks, especially easy, well-practiced ones, have low perceptual loads; these **low-load tasks** use up only a small amount of the person's processing capacity. Tasks that are difficult and perhaps not as well practiced are **high-load tasks** and use more of a person's processing capacity.

Sophie Forster and Lavie (2008) studied the role of processing capacity and perceptual load in determining distraction by presenting displays like the one in **Figure 4.7a**. The participants' task was to respond as quickly as possible when they identified a target, either X or N. Participants pressed one key if they saw the X and another key if they saw the N. This task is easy for displays like the one on the left in Figure 4.7a, in which the target is surrounded by just one type of letter, like the small o's. However, the task becomes harder when the target is surrounded by different letters, as in the display on the right. This difference is reflected in the reaction times, with the hard task resulting in longer reaction times than the easy task. However, when a task-irrelevant stimulus—like the unrelated cartoon character shown in **Figure 4.7b**—is flashed below the display, responding slows for the easy task more than for the hard task.

Lavie explains results such as the ones in Figure 4.7b in terms of her **load theory of attention**, diagrammed in **Figure 4.8**, in which the circle represents the person's processing capacity and the shading represents the portion that is used up by a task. **Figure 4.8a** shows that with the low-load task, there is still processing capacity left. This means that resources are available to process the task-irrelevant stimulus (like the cartoon character), and even

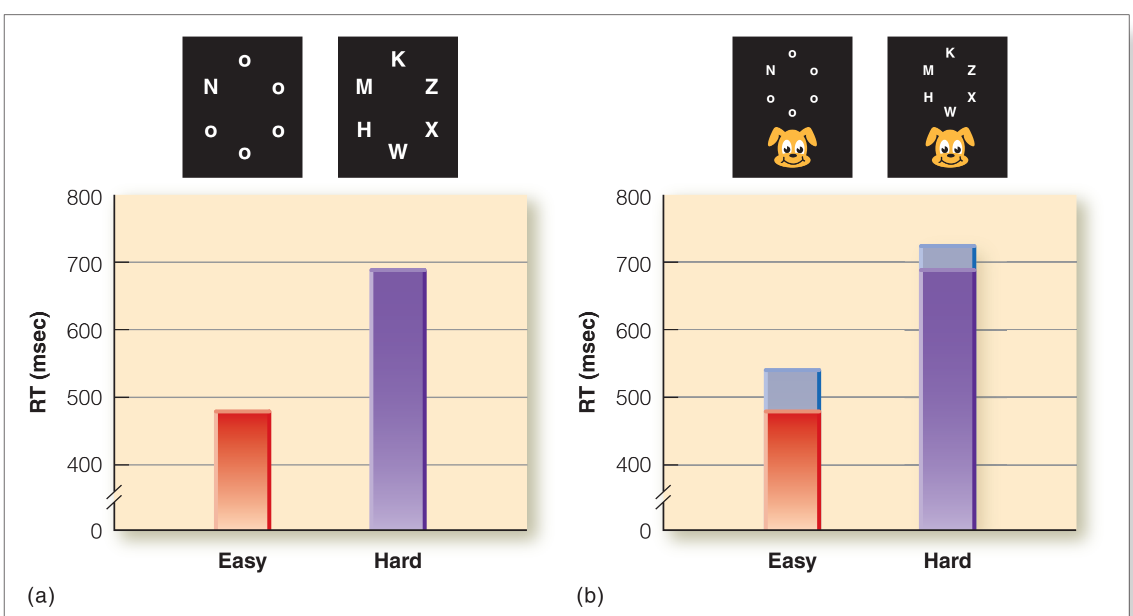

Figure 4.7 The task in Forster and Lavie's (2008) experiment was to indicate the identity of a target (X or N) as quickly as possible in displays like the ones shown here. (a) The reaction time for the easy condition like the display on the left, in which the target is accompanied by small o's, is faster than the reaction time for the hard condition, in which the target is accompanied by other letters. (b) Flashing a distracting cartoon character near the display increases the reaction time for the easy task more than it does for the hard task. The increase for each task is indicated by the gray extensions of the bars.

(Source: Adapted from S. Forster & N. Lavie, Failures to ignore entirely irrelevant distractors: The role of load, *Journal of Experimental Psychology: Applied*, 14 , 73–83, 2008.)

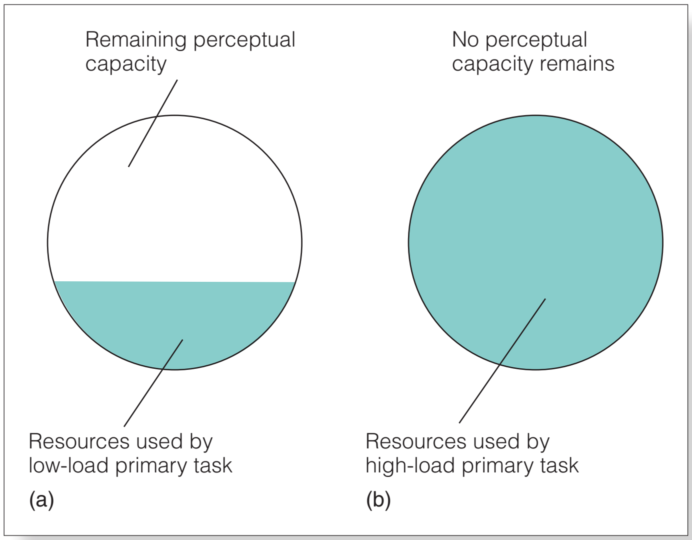

➤ **Figure 4.8** The load theory of attention: (a) Low-load tasks that use few cognitive resources may leave resources available for processing unattended task-irrelevant stimuli, whereas (b) high-load tasks that use all of a person's cognitive resources don't leave any resources to process unattended task-irrelevant stimuli.

though the person was told not to pay attention to the task-irrelevant stimulus, it gets processed and slows down responding.

Figure 4.8b shows a situation in which all of a person's processing capacity is being used for a high-load task, such as the hard task in the experiment. When this occurs, no resources remain to process other stimuli, so irrelevant stimuli can't be processed and they have little effect on performance of the task. Thus, if you are carrying out a hard, high-load task, no processing capacity remains, and you are less likely to be distracted (as Roger found when he was focusing on the hard math problems). However, if you are carrying out an easy, low-load task, the processing capacity that remains is available to process task-irrelevant stimuli (as Roger found out when he was distracted from his easy cell phone game).

The ability to ignore task-irrelevant stimuli is a function not only of the load of the task you are trying to do but also of how powerful the task-irrelevant stimulus is. For example, while Roger was able to ignore the conversation in the library while he was focused on the difficult math problems, a loud siren, indicating fire, would probably attract his attention. An example of a situation in which task-irrelevant stimuli are difficult to ignore is provided by the *Stroop effect*, described in the following demonstration.

DEMONSTRATION The Stroop Effect

Look at **Figure 4.9**. Your task is to name, as quickly as possible, the color of ink used to print each of the shapes. For example, starting in the upper-left corner and going across, you would say, "red, blue, . . ." and so on. Time yourself (or a friend you have enlisted to do this task), and determine how many seconds it takes to report the colors of all the shapes. Then repeat the same task for **Figure 4.10**, remembering that your task is to specify the color of the ink, not the color name that is spelled out.

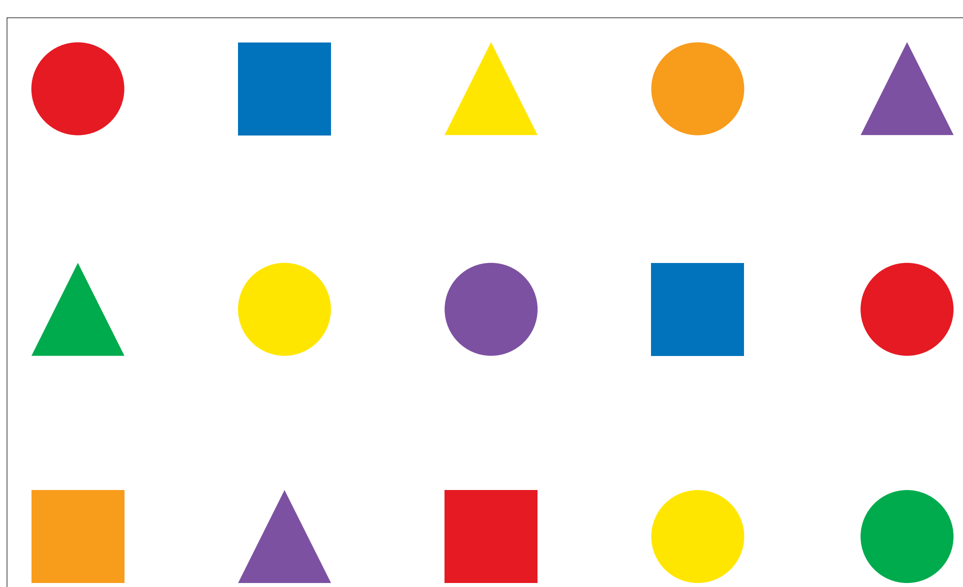

➤ **Figure 4.9** Name the color of the ink used to print these shapes.

YELLOW	RED	BLUE	PURPLE	GREEN
ORANGE	YELLOW	GREEN	BLUE	RED
GREEN	PURPLE	ORANGE	RED	BLUE

➤ **Figure 4.10** Name the color of the ink used to print these words.

If you found it harder to name the colors of the words than the colors of the shapes, then you were experiencing the **Stroop effect**, which was first described by J. R. Stroop in 1935. This effect occurs because the names of the words cause a competing response and therefore slow responding to the target—the color of the ink. In the Stroop effect, the task-irrelevant stimuli are extremely powerful, because reading words is highly practiced and has become so automatic that it is difficult not to read them (Stroop, 1935).

The approaches to attention we have described so far—early information processing models and Lavie's load approach—are concerned with the ability to focus attention on a particular image or task. But in everyday experience you often shift your attention from place to place, either by moving your eyes or by shifting attention "in your mind" without moving your eyes. We discuss such shifts in attention next.

TEST YOURSELF 4.1

1. Give examples of situations that illustrate the following: selective attention, distraction, divided attention, attentional capture, and scanning.
2. How was the dichotic listening procedure used to determine how well people can focus on the attended message and how much information can be taken in from the unattended message? What is the cocktail party effect, and what does it demonstrate?
3. Describe Broadbent's model of selective attention. Why is it called an early selection model?
4. What were the results of experiments by Moray (words in the unattended ear) and Gray and Wedderburn ("Dear Aunt Jane")? Why are the results of these experiments difficult to explain based on Broadbent's filter model of attention?
5. Describe Treisman's attenuation model. First indicate why she proposed the theory, then how she modified Broadbent's model to explain some results that Broadbent's model couldn't explain.
6. Describe MacKay's "bank" experiment. Why does his result provide evidence for late selection?
7. Describe the Forster and Lavie experiment on how processing capacity and perceptual load determine distraction. What is the load theory of attention?
8. What is the Stroop effect? What does it illustrate about task-irrelevant stimuli?

Directing Attention by Scanning a Scene

Attention, according William James, involves "withdrawing from some things in order to effectively deal with others." Think about what this means when applied to everyday situations. There are lots of "things" out there that are potential objects of our attention, but we attend to some things and ignore others. How do we accomplish this, and how does directing our attention affect our experience? We begin by considering how we can direct our attention by moving our eyes to look at one place after another.

Scanning a Scene With Eye Movements

See how many people you can identify in **Figure 4.11** in a minute. Go!

As you did this task, you probably noticed that you had to scan the scene, checking each face in turn, in order to identify it. Scanning is necessary because good detail vision occurs only for things you are looking at directly.

Another way to experience the fact that we have to look directly at things we want to see in detail is to look at the word at the end of this line and, without moving your eyes, see how many words you can read to the left. If you do this without cheating (resist the urge to look to the left!), you will find that although you can read the word you are looking at, you can read only a few of the words that are farther off to the side.

Both of these tasks illustrate the difference between central vision and peripheral vision. *Central vision* is the area you are looking at. *Peripheral vision* is everything off to the side. Because of the way the retina is constructed, objects in central vision fall on a small area called the *fovea*, which has much better detail vision than the peripheral retina, on

Kevin Mazur/WireImage/Getty Images

➤ **Figure 4.11** How many people can you identify in this photo in 1 minute?

which the rest of the scene falls. Thus, as you scanned the scene in Figure 4.11, you were aiming your fovea at one face after another. Each time you briefly paused on one face, you were making a **fixation**. When you moved your eye to observe another face, you were making a **saccadic eye movement**—a rapid, jerky movement from one fixation to the next.

It isn't surprising that you were moving your eyes from one place to another, because you were trying to identify as many people as possible. But it may surprise you to know that even when you are freely viewing an object or scene without searching for anything in particular, you move your eyes about three times per second. This rapid scanning is shown in **Figure 4.12**, which is a pattern of fixations (dots) separated by saccadic eye movements (lines) that occurred as a participant viewed the picture of the fountain. Shifting attention from one place to another by moving the eyes is called **overt attention** because we can see attentional shifts by observing where the eyes are looking.

Maria Wachala/Getty Images (Scanning measurements by James Brockmole)

➤ **Figure 4.12** Scan path of a person freely viewing a picture. Fixations are indicated by the yellow dots and eye movements by the red lines. Notice that this person looked preferentially at areas of the picture such as the statues but ignored areas such as the water, rocks, and buildings.

We will now consider two factors that determine how people shift their attention by moving their eyes: bottom-up, based primarily on physical characteristics of the stimulus; and top-down, based on cognitive factors such as the observer's knowledge about scenes and past experiences with specific stimuli.

Scanning Based on Stimulus Salience

Attention can be influenced by **stimulus salience**—the physical properties of the stimulus, such as color, contrast, or movement. Capturing attention by stimulus salience is a bottom-up process because it depends solely on the pattern of light and dark, color and contrast in a stimulus (Ptak, 2012). For example, the task of finding the people with blonde hair in Figure 4.11 would involve bottom-up processing because it involves responding to the physical property of color, without considering the meaning of the image (Parkhurst et al., 2002). Determining how salience influences the way we scan a scene typically involves analyzing characteristics such as color, orientation, and intensity at each location in the scene and then combining these values to create a **saliency map** of the scene (Itti & Koch, 2000; Parkhurst et al., 2002; Torralba et al., 2006). For example, the person in red in **Figure 4.13** would get high marks for salience, both for the brightness of the color and because it contrasts with the expanse of white, which has lower salience because it is homogeneous.

Experiments in which people's eyes were tracked as they observed pictures have found that the first few fixations are more likely on high-salience areas. But after the first few fixations, scanning begins to be influenced by top-down, or cognitive, processes

Ales Fevzer/Corbis Documentary/Getty Images

➤ **Figure 4.13** The red shirt is visually salient because it is bright and contrasts with its surroundings.

that depend on things such as the observers' goals and expectations determined by their past experiences in observing the environment (Parkhurst et al., 2002).

Scanning Based on Cognitive Factors

One way to show that where we look isn't determined only by stimulus salience is by checking the eye movements of the participant looking at the scene in Figure 4.12. Notice that the person never looks at the bright blue water even though it is salient due to its brightness, color, and position near the front of the scene. The person also ignored the rocks and columns and several other prominent architectural features. Instead, the person focused on aspects of the fountain that might be more interesting, such as the statues. It is important to note, however, that just because this person looked at the statues doesn't mean everyone would. Just as there are large variations between people, there are variations in how people scan scenes (Castelhano & Henderson, 2008; Noton & Stark, 1971). Thus, another person, who might be interested in the architecture of the buildings, might look less at the statues and more at the building's windows and columns.

This example illustrates top-down processing, because scanning is influenced by preferences a person brings to the situation. Top-down processing also comes into play when scanning is influenced by **scene schemas**—an observer's knowledge about what is contained in typical scenes (see regularities of the environment, page 74). Thus, when Melissa Võ and John Henderson (2009) showed pictures like the ones in **Figure 4.14**, observers looked longer at the printer in **Figure 4.14b** than the pot in **Figure 4.14a** because a printer is less likely to be found in a kitchen. People look longer at things that seem out of place in a scene because their attention is affected by their knowledge of what is usually found in the scene.

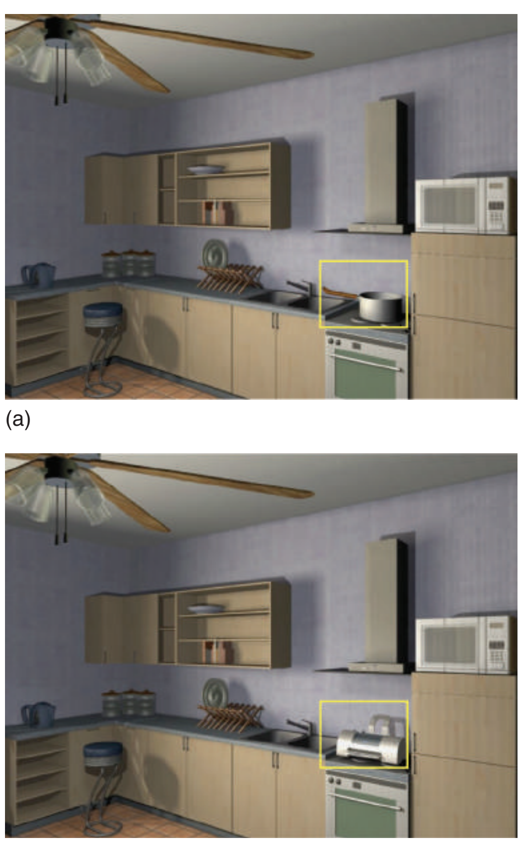

Figure 4.14 Stimuli used by Võ and Henderson (2009). Observers spent more time looking at the printer in (b) than at the pot in (a), shown inside the yellow rectangles (which were not visible to the observers).

(Source: M.L.-H.Vo, & J. M. Henderson, Does gravity matter? Effects of semantic and syntactic inconsistencies on the allocation of attention during scene perception, *Journal of Vision*, 9, 3:24, 1–15, Figure 1)

Another example of how cognitive factors based on knowledge of the environment influences scanning is an experiment by Hiroyuki Shinoda and coworkers (2001) in which they measured observers' fixations and tested their ability to detect traffic signs as they drove through a computer-generated environment in a driving simulator. They found that the observers were more likely to detect stop signs positioned at intersections than those positioned in the middle of a block, and that 45 percent of the observers' fixations occurred close to intersections. In this example, the observers are using learning about regularities in the environment (stop signs are usually at corners) to determine when and where to look for stop signs.

Scanning Based on Task Demands

The examples in the last section demonstrate that knowledge of various characteristics of the environment can influence how people direct their attention. However, the last example, in which participants drove through a computer-generated environment, was different from the rest. The difference is that instead of looking at pictures of stationary scenes, participants were interacting with the environment. This kind of situation, in which people are shifting their attention from one place to another as they are doing things, occurs when people are moving through the environment, as in the driving example, and when people are carrying out specific tasks.

Because many tasks require attention to different places as the task unfolds, it isn't surprising that the timing of when people look at specific places is determined by the sequence of actions involved in the task. Consider, for example, the pattern of eye movements in **Figure 4.15**, which were measured as a person was making a peanut butter sandwich. The process of making the sandwich begins with the movement of a slice of bread from the bag (A) to the plate (B). Notice that this operation is accompanied by an eye movement from the bag to the plate. The observer then looks at the peanut butter jar just before it is lifted and looks at the top just before it is removed (C). Attention then shifts to the knife, which is picked up and used to scoop the peanut butter and spread it on the bread (Land & Hayhoe, 2001).

The key finding of these measurements, and also of another experiment in which eye movements were measured as a person prepared tea (Land et al., 1999), is that the person's eye movements were determined primarily by the task. The person fixated on few objects or areas that were irrelevant to the task, and eye movements and fixations were closely linked to the action the person was about to take. Furthermore, the eye movement usually preceded a motor action by a fraction of a second, as when the person first fixated on the peanut butter jar and then reached over to pick it up. This is an example of the "just in time" strategy—eye movements occur just before we need the information they will provide (Hayhoe & Ballard, 2005; Tatler et al., 2011).

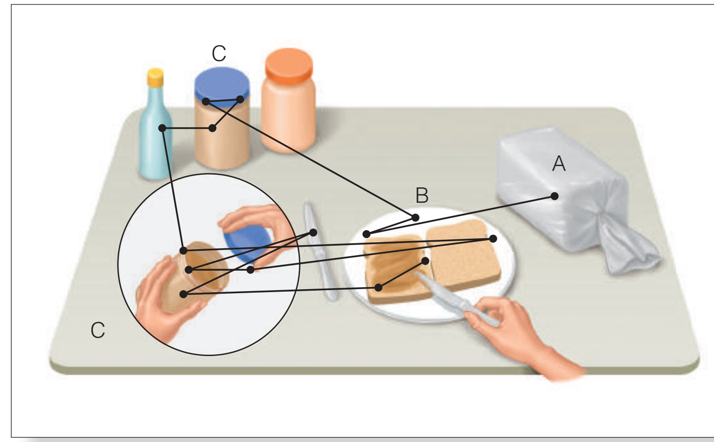

➤ **Figure 4.15** Sequence of fixations of a person making a peanut butter sandwich. The first fixation is on the loaf of bread.

(Source: Adapted from M. F. Land, N. Mennie, & J. Rusted, The roles of vision and eye movements in the control of activities of daily living, *Perception*, 28, 11, 1311–1328. Copyright © 1999 by Pion Ltd, London. Reproduced by permission. www.pion.co.uk and www.envplan.com.)

The examples we have described in connection with scanning based on cognitive factors and task demands have something in common: They all provide evidence that scanning is influenced by people's *predictions* (Henderson, 2017). Scanning anticipates what a person is going to do next as they make a peanut butter and jelly sandwich; scanning anticipates that stop signs are most likely to be located at intersections; and pausing scanning to look longer at an unexpected object occurs when a person's expectations are violated, as when a printer unexpectedly appears in a kitchen.

Outcomes of Attention

What do we gain by attending? Based on the last section, which described *overt attention* that is associated with eye movements, we might answer that question by stating that shifting attention by moving our eyes enables us to see places of interest more clearly. This is extremely important, because it places the things we're interested in front-and-center where they are easy to see.

But some researchers have approached attention not by measuring factors that influence eye movements, but by considering what happens when we shift our attention without making eye movements. Shifting attention while keeping the eyes still is called **covert attention**, because the attentional shift can't be seen by observing the person. This type of attending involves shifting attention "with the mind" as you might do when you are paying attention to something off to the side while still looking straight ahead. (This has also been described as "looking out of the corner of your eye.")

One reason some researchers have studied covert attention is that it is a way of studying what is happening in the mind without the interference of eye movements. We will now consider research on covert attention, which shows how shifting attention "in the mind" can affect how quickly we can respond to locations and to objects, and how we perceive objects.

Attention Improves Our Ability to Respond to a Location

In a classic series of studies, Michael Posner and coworkers (1978) asked whether paying attention to a location improves a person's ability to respond to stimuli presented there. To answer this question, Posner used a procedure called **precueing**.

METHOD Precueing

The general principle behind a precueing experiment is to determine whether presenting a cue indicating where a test stimulus will appear enhances the processing of the target stimulus. The participants in Posner and coworkers' (1978) experiment kept their eyes stationary throughout the experiment, always looking at the + in the display in **Figure 4.16**, so Posner was measuring covert attention.

Participants first saw an arrow cue (as shown in the left panel) indicating on which side of the display they should focus their attention. In **Figure 4.16a**, the arrow cue indicates that they should focus their attention to the right (while looking steadily at the +). Their task was to press a key as rapidly as possible when a target square was presented off to the side (as shown in the right panel). The trial shown in **Figure 4.16a** is a valid trial because the target square appears on the side indicated by the cue arrow. On 80 percent of the trials, the cue arrow directed participants' attention to the side where the target square appeared. However, on 20 percent of the trials, the arrow directed the participant's attention away from where the target was to appear (**Figure 4.16b**). These were the invalid trials. On both the valid and invalid trials, the participant's task was the same—to press the key as quickly as possible when the target square appeared.

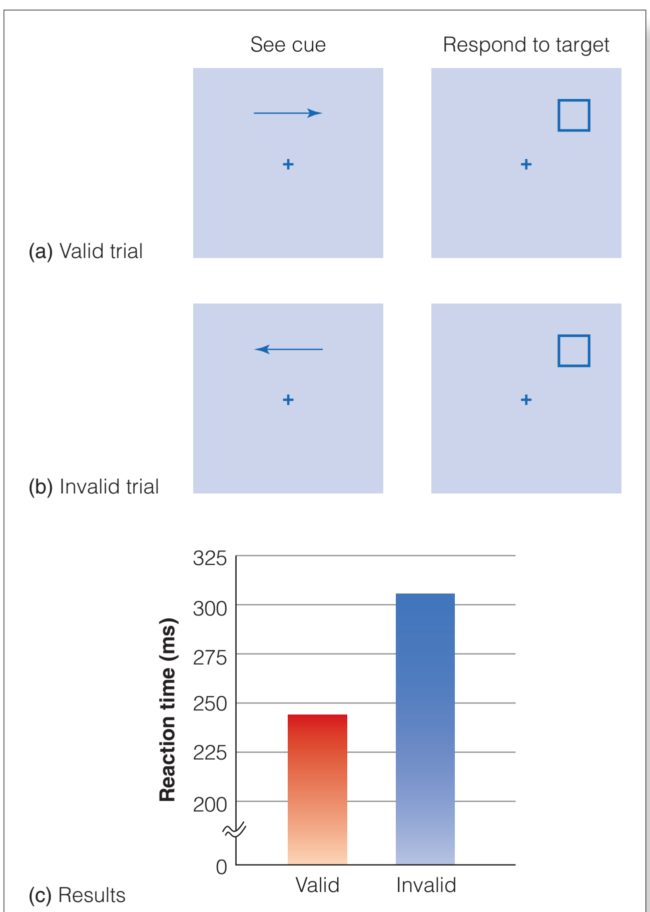

Figure 4.16 Procedure for (a) valid trials and (b) invalid trials in Posner et al.'s (1978) precueing experiment; (c) the results of the experiment. The average reaction time was 245 ms for valid trials but 305 ms for invalid trials.

(Source: M. I. Posner, M. J. Nissen, & W. C. Ogden, Modes of perceiving and processing information. Copyright © 1978 by Taylor & Francis Group LLC–Books.)

The results of this experiment, shown in **Figure 4.16c**, indicate that participants reacted to the square more rapidly when their attention was focused on the location where the signal was to appear. Posner interpreted this result as showing that information processing is more effective at the place where attention is directed. This result and others like it gave rise to the idea that attention is like a spotlight or zoom lens that improves processing when directed toward a particular location (Marino & Scholl, 2005).

Attention Improves Our Ability to Respond to Objects

In addition to covertly attending to locations, as in Posner's experiment, we can also covertly attend to specific objects. We will now consider some experiments that show that (1) attention can enhance our response to objects and (2) when attention is directed to one place on an object, the enhancing effect of that attention spreads to other places on the object.

Consider, for example, the experiment diagrammed in **Figure 4.17** (Egly et al., 1994). As participants kept their eyes on the +, one end of the rectangle was briefly highlighted (**Figure 4.17a**). This was the cue signal that indicated where a target, a dark square (**Figure 4.17b**), would probably appear. In this example, the cue indicates that the target is likely to appear in position A, at the upper part of the right rectangle, and the target is, in fact, presented at A. (The letters used to illustrate positions in our description did not appear in the actual experiment.)

The participants' task was to press a button when the target was presented anywhere on the display. The numbers indicate the reaction times, in milliseconds, for three target locations when the cue signal had been presented at A. Not surprisingly, participants responded most rapidly when the target was presented at A, where the cue had been presented. However, the most interesting result is that participants responded more rapidly when the target was presented at B (reaction time = 358 ms) than when the target was presented at C (reaction time = 374 ms). Why does this occur? It can't be because B is closer to A than C, because B and C are exactly the same distance from A. Rather, B's advantage occurs because it is located within the object that was receiving the participant's attention. Attending at A, where the cue was presented, causes the maximum effect at A, but the effect of this attention spreads throughout the object so some enhancement occurs at B as well. The faster responding that occurs when enhancement spreads within an object is called the **same-object advantage** (Marino & Scholl, 2005; also see Driver & Baylis, 1989, 1998; Katzner et al., 2009; Lavie & Driver, 1996; and Malcolm & Shomstein, 2015 for more demonstrations of how attention spreads throughout objects).

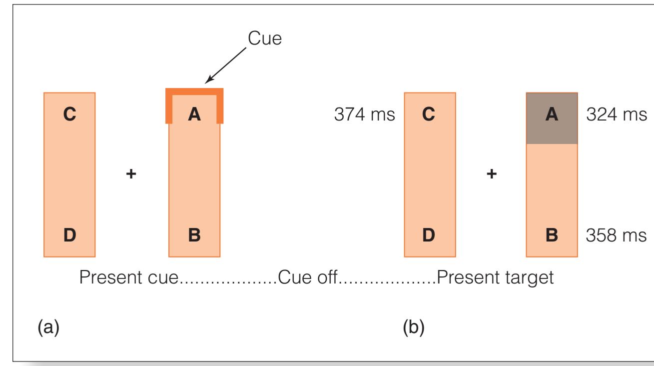

➤ **Figure 4.17** In Egly and coworkers' (1994) experiment, (a) a cue signal appears at one place on the display, then the cue is turned off and (b) a target is flashed at one of four possible locations, A, B, C, or D. The participants' task was to press a button when the target was presented anywhere on the display. Numbers are reaction times in ms for positions A, B, and C when the cue signal appeared at position A.

Attention Affects Perception

Returning to the quote by James at the beginning of the chapter, let's focus on his description of attention to objects as "taking possession by the mind *in clear and vivid form*." The phrase *in clear and vivid form* suggests that attending to an object makes it more clear and vivid—that is, attention affects perception. More than 100 years after James's suggestion, many experiments have shown that attended objects are perceived to be bigger and faster, and to be more richly colored and have better contrast than non attended objects (Anton-Erxleben et al., 2009; Carrasco et al., 2004; Fuller & Carrasco, 2006; Turatto et al., 2007). Attention therefore not only causes us to respond faster to locations and objects but affects how we perceive the object (Carrasco, 2011).

Attention Affects Physiological Responding

Attention has a number of different effects on the brain. One effect is to increase activity in areas of the brain that represent the attended location.

Attention to Locations Increases Activity in Specific Areas of the Brain What happens in the brain when people shift their attention to different locations while keeping their eyes stationary? Ritobrato Datta and Edgar DeYoe (2009) answered this question by measuring brain activity using fMRI as participants kept their eyes fixed on the center of the display in **Figure 4.18** and shifted their attention to different locations in the display.

The colors in the circles in **Figure 4.18b** indicate the area of brain that was activated when the participant directed his or her attention to different locations indicated by the letters on the stimulus in Figure 4.18a. Notice that the yellow "hot spot," which is the place of greatest activation, moves out from the center and also becomes larger as attention

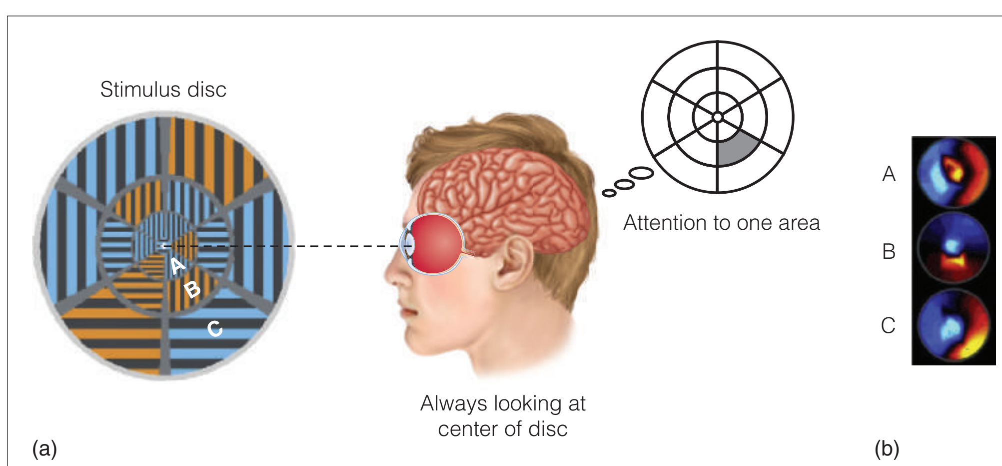

➤ **Figure 4.18** (a) Participants in Datta and DeYoe's (2009) experiment directed their attention to different areas of this circular display while keeping their eyes fixed on the center of the display. (b) Activation of the brain that occurred when participants attended to the areas indicated by the letters on the stimulus disc. The center of each circle is the place on the brain that corresponds to the center of the stimulus. The yellow "hot spot" is the area of the brain that is maximally activated by attention.

(Source: From R. Datta & E. A. DeYoe, I know where you are secretly attending! The topography of human visual attention revealed with fMRI, *Vision Research*, 49, 1037–1044, 2009.)

is directed farther from the center. By collecting brain activation data for all of the locations on the stimulus, Datta and DeYoe created "attention maps" that show how directing attention to a specific area of space activates a specific area of the brain.

What makes this experiment even more interesting is that after attention maps were determined for a particular participant, that participant was told to direct his or her attention to a "secret" place, unknown to the experimenters. Based on the location of the resulting yellow "hot spot" in the brain, the experimenters were able to predict, with 100 percent accuracy, the "secret" place where the participant was attending.

Attention Changes the Representation of Objects Across the Cortex Datta and DeYoe's "hot spot" experiment is an elegant demonstration of how attention directed to a specific location results in enhanced activity at one place in the cortex. But what about a situation in which people might be directing their attention to numerous different locations as they search for something in a naturalistic environment? Tolga Cukur and coworkers (2013) considered this question by determining how attention affects the way different types of objects are represented across the brain as a whole.

The starting point for Cukur's experiment was Alex Huth's (2012) brain map that we described in Chapter 2 (see Figure 2.20). Huth's map illustrates how different categories of objects and actions are represented by activity that is distributed across a large area of the brain. Huth determined this map by having participants view movies in a scanner, and using fMRI to determine brain activity when different things were happening on the screen (see Figure 2.19).

Cukur did the same thing as Huth (they were working in the same laboratory and were involved in both papers), but instead of having his observers passively view the movies, he gave them a task that involved searching for either "humans" or "vehicles." A third group passively viewed the films, as in Huth's experiment. **Figure 4.19** shows what happened by plotting how a single voxel in the brain (see page 41 to review *voxel*) responded to different types of stimuli under two different search conditions. Notice in (a) that when the observer is searching for "humans" in the movie, the voxel responds well to "person," slightly to "animal," and hardly at all to "building" and "vehicle." However, in (b), when the observer is searching for "vehicle," the voxel's tuning shifts so it now responds well to "vehicle," slightly to "building," but not to "person" or "animal."

By analyzing the data from tens of thousands of voxels across the brain, Cukur created the whole brain maps shown in **Figure 4.20**. The colors indicate tuning to different categories. The most obvious difference between the search-for-people brain and the search-for-vehicles brain occurs at the top of the brain in this view. Notice that in the person condition there are more yellows and greens, which represent people or things related to people like body parts, animals, groups, and talking. However, in the vehicles condition,

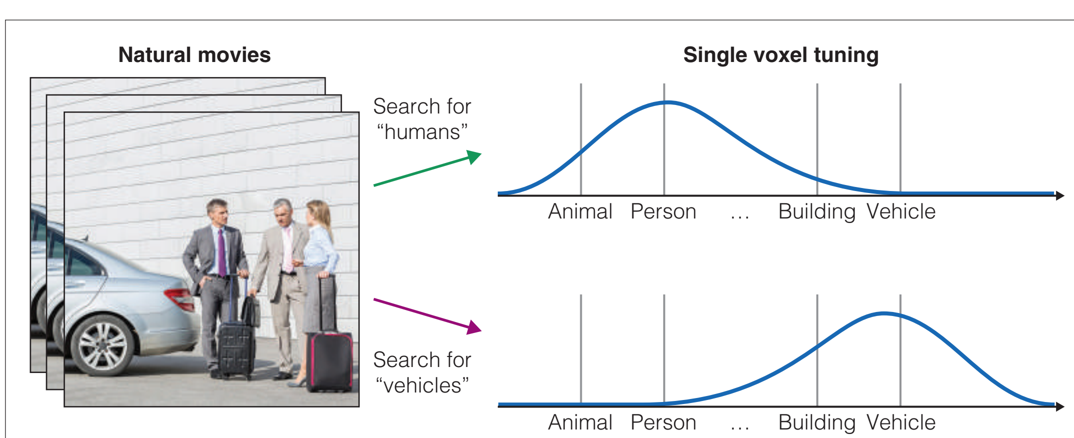

moodboard/Alamy Stock Photo

Figure 4.19 How tuning of a single voxel is affected by attention to (a) humans and (b) vehicles.

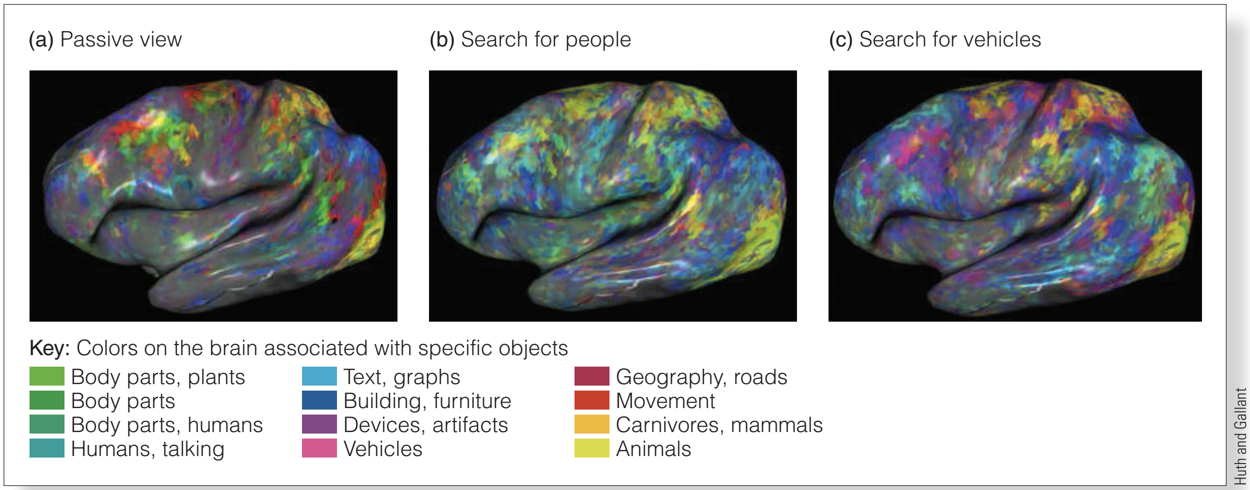

➤ **Figure 4.20** The map of categories on the brain changes for viewing a film. Colors indicate activation caused by different categories of stimuli. (a) Passive view indicates activation when not searching for anything. (b) Search for people causes activation indicated by yellow and green, which stand for people and things related to people. (c) Search for vehicles causes activation indicated by reds, which stand for vehicles and things related to vehicles.

colors shift to reds, which represent vehicles or things related to vehicles such as movement, road, and devices.

An important feature of these brain maps is that looking for a particular category shifts responding to the category and to additional things related to that category, so looking for people also affects responding to groups and clothing. Cukur calls this effect **attentional warping**—the map of categories on the brain changes so more space is allotted to categories that are being searched for, and this effect occurs even when the attended category isn't present in the movie. For example, when a person is on the lookout for vehicles, the brain becomes "warped" or "tuned" so that large areas respond best to vehicles and things related to vehicles. Then, when a vehicle, a road, or movement appears in a scene, a large response occurs. Other things, which the person is not looking for at the moment, would cause smaller responses.

TEST YOURSELF 4.2

1. What is the difference between central vision and peripheral vision? How is this difference related to overt attention, fixations, and eye movements?
2. What is stimulus salience? How is it related to attention?
3. Describe some examples of how attention is determined by cognitive factors. What is the role of scene schemas?
4. Describe the peanut butter experiment. What does the result tell us about the relation between task demands and attention?
5. What is covert attention? Describe the precueing procedure used by Posner. What does the result of Posner's experiment indicate about the effect of attention on information processing?
6. Describe the Egly precueing experiment. What is the same-object advantage, and how was it demonstrated by Egly's experiment?
7. What are three behaviorally measured outcomes of attention?
8. Describe how Data and DeYoe showed that attention to a location affects activity in the brain.
9. Describe Cukur's experiment, which showed how attention changes the representation of objects across the cortex.

Divided Attention: Can We Attend to More Than One Thing at a Time?

Our emphasis so far has been on attention as a mechanism for focusing on one thing at a time. We have seen that sometimes we take in information from a task-irrelevant stimulus, even when we are trying ignore irrelevant stimuli, as in Forster and Lavie's experiment and the Stroop task. But what if you want to purposely distribute your attention among a few tasks? Is it possible to pay attention to more than one thing at a time? Although you might be tempted to answer "no," based on the difficulty of listening to two conversations at once, there are many situations in which divided attention—the distribution of attention among two or more tasks—can occur, as when Roger was able to play his cell phone game and listen in on the nearby conversation. Also, people can simultaneously drive, have conversations, listen to music, and think about what they're going to be doing later that day (although this may not hold for difficult driving conditions). As we will see, the ability to divide attention depends on a number of factors, including practice and the difficulty of the task.

Divided Attention Can Be Achieved With Practice: Automatic Processing

Experiments by Walter Schneider and Richard Shiffrin (1977) involved divided attention because they required the participant to carry out two tasks simultaneously: (1) holding information about target stimuli in memory and (2) paying attention to a series of "distractor" stimuli to determine whether one of the target stimuli is present among these distractor stimuli. **Figure 4.21** illustrates the procedure. The participant was shown a memory set like the one in **Figure 4.21a**, consisting of one to four characters called target stimuli. The memory set was followed by rapid presentation of 20 "test frames," each of which contained distractors (**Figure 4.21b**). On half of the trials, one of the frames contained a target stimulus from the memory set. A new memory set was presented on each trial, so the targets changed from trial to trial, followed by new test frames. In this example, there is one target stimulus in the memory set, there are four stimuli in each frame, and the target stimulus 3 appears in one of the frames.

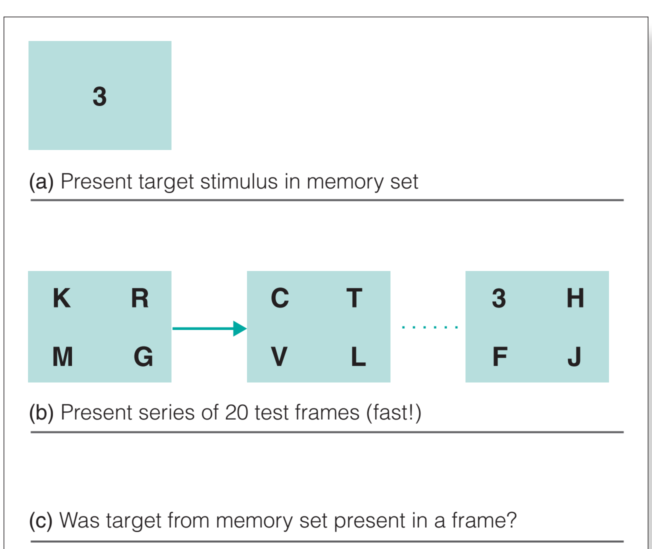

Figure 4.21 Sample stimuli for Schneider and Shiffrin's (1977) experiment. In this experiment, there is one target stimulus in the memory set (the 3) and four stimuli in each frame. The target appears in the last frame in this example.

(Source: R. M. Shiffrin & W. Schneider, Controlled and automatic human information processing: Perceptual learning, automatic attending, and a general theory, *Psychological review*, 84, 127–190, 1977.)

At the beginning of the experiment, the participants' performance was only 55 percent correct; it took 900 trials for performance to reach 90 percent (**Figure 4.22**). Participants reported that for the first 600 trials, they had to keep repeating the target items in each memory set in order to remember them. (Although targets were always numbers and distractors letters, remember that the actual targets and distractors changed from trial to trial.) However, participants reported that after about 600 trials, the task had become automatic: The frames appeared and participants responded without consciously thinking about it. They would do this even when as many as four targets had been presented.

What this means, according to Schneider and Shiffrin, is that practice made it possible for participants to divide their attention to deal with all of the target and test items simultaneously. Furthermore, the many trials of practice resulted in **automatic processing**, a type of processing that occurs (1) without intention (it happens automatically without the person intending to do it) and (2) at a cost of only some of a person's cognitive resources.

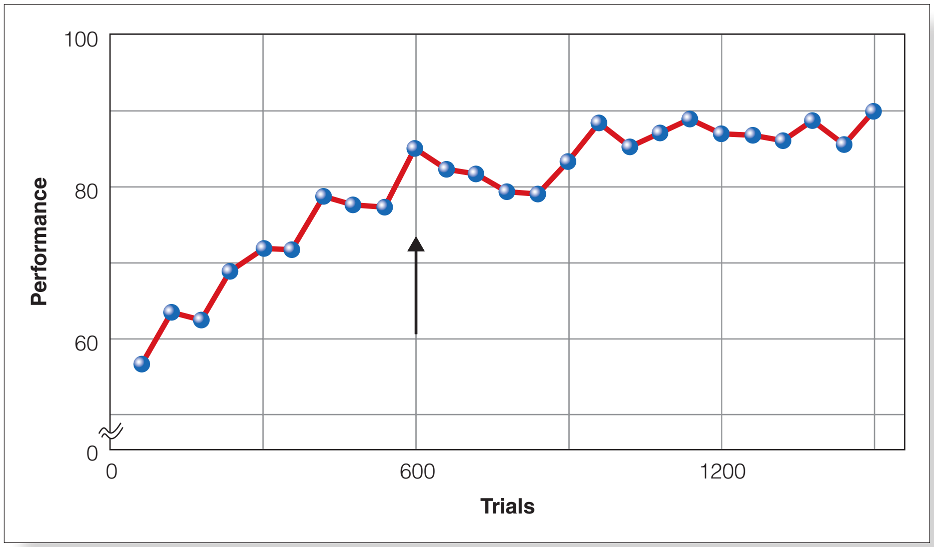

➤ **Figure 4.22** Improvement in performance with practice in Schneider and Shiffrin's (1977) experiment. The arrow indicates the point at which participants reported that the task had become automatic. This is the result of experiments in which there were four target stimuli in the memory set and two stimuli in each test frame

(Source: R. M. Shiffrin & W. Schneider, Controlled and automatic human information processing: Perceptual learning, automatic attending, and a general theory, *Psychological review*, 84, 127–190, 1977.)

Real-life experiences are filled with examples of automatic processing because there are many things that we have been practicing for years. For example, have you ever wondered, after leaving home, whether you had locked the door and then returned to find that you had? Locking the door has, for many people, become such an automatic response that they do it without paying attention. Another example of automatic processing (which is sometimes scary) occurs when you have driven somewhere and can't remember the trip once you get to your destination. In many cases, this involves being "lost in thought" about something else, yet driving has become so automatic that it seems to take care of itself (at least until a traffic "situation" occurs, such as road construction or another car cutting in front of you). Finally, you may carry out many motor skills, such as touch-typing or texting, automatically, without attention. Try paying attention to what your fingers are doing while typing and notice what happens to your performance. Concert pianists have reported that if they start paying attention to their fingers while they are playing, their performance falls apart.

Divided Attention Becomes More Difficult When Tasks Are Harder

What Schneider and Shiffrin's experiment shows is that divided attention is possible for some well-practiced tasks. However, in other experiments, they found that if task difficulty is increased—by using letters for both targets and distractors and by changing targets and distractors on each trial so a target on one trial can be a distractor on another—then automatic processing is not possible even with practice (also see Schneider & Chein, 2003).

An example of divided attention becoming difficult when the task is made too hard is provided by driving. You may find it easy to drive and talk at the same time if traffic is light on a familiar road. But if traffic increases, you see a flashing "Construction Ahead" sign, and the road suddenly becomes rutted, you might have to stop a conversation or turn off the radio, so you can devote all of your cognitive resources to driving. Because of the importance of driving in our society and the recent phenomenon of people talking on cell phones and texting while driving, researchers have begun to investigate the consequences of attempting to divide attention between driving and other distracting activities.

Distractions

The environment is full of distractions—things that direct our attention away from something we are doing. A source of distractions that has become widespread in the last few decades is cell phones, tablets, and computers, and one of the most dangerous consequences of this source of distraction occurs while driving.

Distractions by Cell Phones while Driving

Driving presents a paradox: in many cases, we are so good at it that we can operate on "autopilot," as when we are driving down a straight highway in light traffic. However, in other cases, driving can become very demanding, as noted earlier, when traffic increases or hazards suddenly present themselves. It is in this latter case that distractions that result in a decrease in attention to driving are particularly dangerous.

The seriousness of driver inattention was verified by a research project called the 100-Car Naturalistic Driving Study (Dingus et al., 2006). In this study, video recorders in 100 vehicles created records of both what the drivers were doing and the view out the front and rear windows. These recordings documented 82 crashes and 771 near crashes in more than 2 million miles of driving. In 80 percent of the crashes and 67 percent of the near crashes, the driver was inattentive in some way 3 seconds beforehand. One man kept glancing down and to the right, apparently sorting through papers in a stop-and-go driving situation, until he slammed into an SUV. A woman eating a hamburger dropped her head below the dashboard just before she hit the car in front of her. One of the most distracting activities was pushing buttons on a cell phone or similar device. More than 22 percent of near crashes involved that kind of distraction, and it is likely that this number may be higher now because of increases in cell phone use since that study.

In a laboratory experiment on the effects of cell phones, David Strayer and William Johnston (2001) gave participants a simulated driving task that required them to apply the brakes as quickly as possible in response to a red light. Doing this task while talking on a cell phone caused participants to miss twice as many of the red lights as when they weren't talking on the phone (**Figure 4.23a**) and also increased the time it took them to apply the brakes (**Figure 4.23b**). Perhaps the most important finding of this experiment is that the same decrease in performance occurred regardless of whether participants used a hands-free or a handheld device.

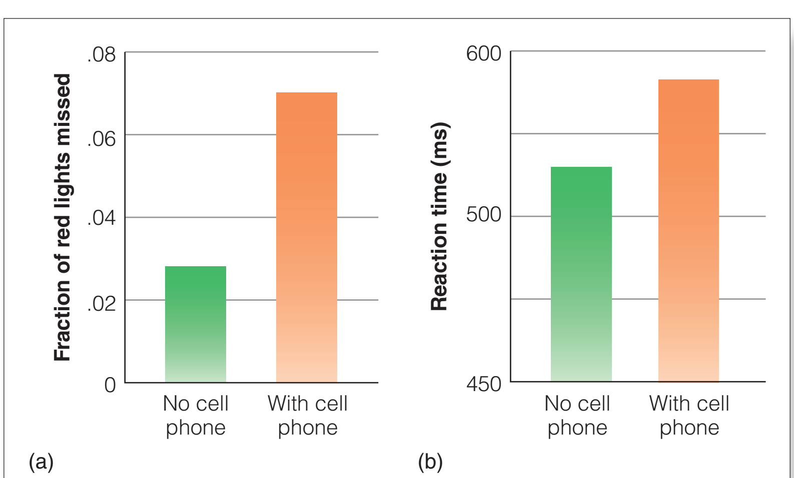

Figure 4.23 Result of Strayer and Johnston's (2001) cell phone experiment. When participants were talking on a cell phone, they (a) missed more red lights and (b) took longer to apply the brakes.

Taking into account results such as these, plus many other experiments on the effects of cell phones on driving, Strayer and coworkers (2013) concluded that talking on the phone uses mental resources that would otherwise be used for driving the car (also see Haigney & Westerman, 2001; Lamble et al., 1999; Spence & Read, 2003; Violanti, 1998). This conclusion that the problem posed by cell phone use during driving is related to the use of mental resources is an important one. The problem isn't driving with one hand. It is driving with fewer mental resources available to focus on driving.

But even though research clearly shows that driving while talking on a cell phone is dangerous, many people believe it doesn't apply to them. For example, in response to a class assignment, one of my students wrote, "I do not believe my driving is affected by talking on the phone. . . . My generation learned to drive when cell phones were already out. I had one before driving, so while learning to drive, I also simultaneously learned to talk on the phone and drive." Thinking such as this may be why 27 percent of adults report that they sometimes text while driving, even in the face of overwhelming evidence that it is dangerous (Seiler, 2015; Wiederhold, 2016). For example, a study by the Virginia Tech Transportation Institute found that truck drivers who send text messages while driving were 23 times more likely to cause a crash or near crash than truckers who were not texting (Olson et al., 2009). Because of results such as these, which indicate that texting is even more dangerous than talking on a cell phone, most states now have laws against text-messaging while driving.

The main message here is that anything that distracts attention can degrade driving performance. And cell phones aren't the only attention-grabbing device found in cars. New car models feature small screens that can display the same apps that are on your phone. Some voice-activated apps enable drivers to make movie or dinner reservations, send and receive texts or emails, and post on Facebook. Ford calls their system an "infotainment system." But a recent study from the AAA Foundation for Traffic Safety, *Measuring Cognitive Distraction in the Automobile*, indicates that perhaps too much information and entertainment isn't a good thing. The study found that voice-activated activities were more distracting, and therefore potentially more dangerous, than either hands-on or hands-free cell phones. The study concludes that "just because a new technology does not take the eyes off the road does not make it safe to be used while the vehicle is in motion" (Strayer et al., 2013).

Distractions by the Internet

There is no question that distracted driving caused by cell phone use impacts the ability to drive safely. But cell phones, and the Internet in general, can also have negative effects on many other aspects of behavior.

Many research studies have documented high usage of cell phones and the Internet. For example, 92 percent of college students report that they have texted, browsed the web, sent pictures, or visited social networks during class time (Tindall & Bohlander, 2012). By checking college students' phone bills (with their permission!), Judith Gold and coworkers (2015) determined that they send an average of 58 text messages a day, and Rosen and coworkers (2013) showed that during a 15-minute study session, students averaged less than 6 minutes on-task before interrupting studying to stretch, watch TV, access websites, or use technology such as texting or Facebook.

Another method of determining ongoing daily behaviors such as texting is *experience sampling*.

METHOD Experience Sampling

Experience sampling was developed to answer the question, "what percentage of the time during the day are people engaged in a specific behavior?" One way this question has been answered is to use a cell phone app that sends people text messages at random times during the day, asking them questions. For example, to determine the frequency of Internet usage, the question might be "Were you on the Internet?" Additional questions could also be inserted, such as "What type of online activity were you engaged in?" with choices such as "social networking," "email," and "browsing." When Moreno and coworkers (2012) sent students text-message probes at random times six times a day, they found that 28 percent of the probes arrived when the student was on the phone or Internet.

How often do you consult your cell phone? If you check your phone constantly, one explanation of your behavior involves **operant conditioning**, a type of learning named by behaviorist B.F. Skinner (1938) (see p. 11), in which behavior is controlled by rewards (called reinforcements) that follow behaviors. A basic principle of operant conditioning is that the best way to ensure that a behavior will continue is to reinforce it intermittently. So when you check your phone for a message and it's not there, well, there's always a chance it will be there the next time. And when it eventually appears, you've been intermittently reinforced, which strengthens future phone-clicking behavior. Some people's dependence on their phone is captured in the following sticker, marketed by Ephemera, Inc: "After a long weekend without your phone, you learn what's really important in life. Your phone." (See Bosker, 2016, for more on how cell phones are programmed to keep you clicking.)

Constant switching from one activity to another has been described as "continuous partial attention" (Rose, 2010), and here is where the problem lies, because as we saw for driving, distraction from a task impairs performance. It isn't surprising, therefore, that people who text more tend to have lower grades (Barks et al., 2011; Kuznekoff et al., 2015; Kuznekoff & Titsworth, 2013; Lister-Landman et al., 2015), and in extreme cases, some people are "addicted" to the Internet, where addiction is defined as occurring when Internet use negatively affects a number of areas of a person's life (for example, social, academic, emotional, and family) (Shek et al., 2016).

What's the solution? According to Steven Pinker (2010), given that the computer and Internet are here to stay, "the solution is not to bemoan technology, but to develop strategies of self-control, as we do with every other temptation in life." This sounds like good advice, but sometimes powerful temptations are difficult to resist. One example, for some people, is chocolate. Another is checking their cell phone. But even if you are able to resist chocolate and your cell phone, there is another distraction that is difficult to resist: the distraction that occurs when your mind wanders.

Distraction Caused by Mind Wandering

Let's return to Roger, who, at the beginning of the chapter, was sitting in the library wondering how to solve some math problems. Even though he is able to ignore the people talking next to him, he suddenly realizes that his mind has drifted away from doing his math problems to thinking about what he's going to do later, and then there's the problem of what to get his girlfriend for her birthday, and then . . . But wait a minute! What happened to the math problem? Roger's mind has been waylaid by **mind wandering**—thoughts coming from within—which have also been called *daydreaming* (Singer, 1975; Smallwood & Schooler, 2015) (**Figure 4.24**).

➤ **Figure 4.24** According to Killingsworth and Gilberts (2010), people are mind wandering about half the time when they are awake. Here, Roger is supposed to be focusing on doing math problems but seems to be drifting off onto other topics.
(Source: Killingsworth and Gilberts, 2010)

One of the properties of mind wandering is that it is extremely prevalent. Matthew Killingsworth and Daniel Gilbert (2010) used the experience sampling technique to contact people at random intervals during the day and ask them "What are you doing right now?" Mind wandering occurred 47 percent of the time and occurred when people were involved in a wide range of activities (**Table 4.1**). So mind wandering is extremely prevalent and, as shown in other studies, is distracting enough to disrupt an ongoing task (Mooneyham & Schooler, 2013). An example of disruption by mind wandering is what happens while reading, when you suddenly realize that you have no idea what you've just read because you were thinking about something else. This phenomenon, called *mindless reading* or *zoned-out reading*, is one example of how mind wandering decreases performance (Smallwood, 2011).

TABLE 4.1
Activities During Which Mind Wandering Occurs, in Order of Frequency.

Most frequent activities are listed first, starting at the top left.

Working	Eating	Playing
Talking/conversing	Reading	Exercising
Using a computer	Shopping, running errands	Walking
Commuting, traveling	Reading	Listening to music
Watching television	Doing housework	Making love
Relaxing	Grooming/self-care	Listening to the radio
Resting/sleeping	Taking care of children	Praying, meditating

Source: From Killingsworth & Gilbert, 2010.

Another property of mind wandering is that it is usually associated with activity in the default mode network (DMN). Remember, from Chapter 2 (page 50), that the DMN becomes activated when a person is not involved in a task. This may seem to contradict the preceding examples, when mind wandering is occurring during tasks like doing math problems or reading. But remember that once a person's mind begins to wander, he or she is no longer focusing their attention on a task. Mind wandering is a big problem if you need to stay focused. However, as we will see later in the book, when we consider memory, problem solving, and creativity, mind wandering also has benefits, such as helping us plan for the future and enhancing creativity.

What Happens When We Don't Attend?

One thing that should be clear from our discussion so far is that attention is a precious, but limited, resource. We can attend to some things but not to everything. Dividing attention is possible but difficult, and there are forces out there in the world trying to distract us from paying attention to what we're supposed to be attending to. (Feel free to take a break here to check your cell phone for messages, but don't be distracted for too long, because there's still more to our story.)

There are many ways of showing that limits exist to our ability to attend, and we can demonstrate this by looking at what happens when we aren't attending to the right place at the right time. If we are paying attention to some things in a scene, we inevitably miss other things. This is dramatically illustrated by a tragic accident that occurred at a swimming pool in Iowa, described as follows by Lyndsey Lanagan-Leitzel and coworkers (2015):

> On 14 July 2010, approximately 175 teenage boys enjoyed a day of swimming at a local pool in Pella, Iowa, as part of a Fellowship of Christian Athletes camp held at Central College. When it was time to board busses to return to their rooms, two boys were discovered missing. A 15-minute search culminated in a tragic discovery—the bodies of two of the boys (ages 14 and 15 years) were discovered motionless on the bottom of the pool. Attempts to revive them failed (Belz, 2010).

What was particularly surprising about these drownings is that although at least 10 lifeguards and 20 camp counselors were observing the swimmers, nobody was aware of their drowning. According to Lanagan-Leitzel, although fatal drownings in lifeguarded pools are rare, there are reasons involving the limits of our ability to attend that might explain why this happened.

Consider the lifeguard's task. They are essentially carrying out a visual scanning task in which their job is to detect a rare event (someone drowning) amidst many similarly appearing distractors (boys splashing in a pool). Apparently, it isn't uncommon for people to drown without excessive splashing, and they often don't yell out for help because they are focusing their energy and attention on trying to breathe. There are other reasons it is sometimes difficult to spot someone who is drowning in a crowded pool, but the message,

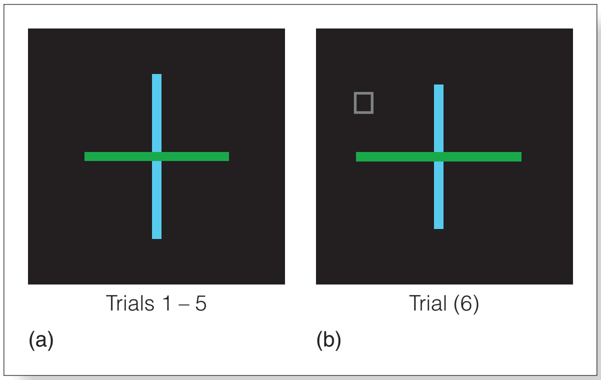

➤ **Figure 4.25** Inattentional blindness experiment. (a) The cross display is presented for five trials. On each trial, one arm of the cross is slightly longer than the other. The participant's task is to indicate which arm (horizontal or vertical) is longer. (b) On the sixth trial, the participants carry out the same task, but a small square or other geometric object is included in the display. After the sixth trial, participants are asked whether they saw anything different than before.

(Source: Adapted from N. Lavie, Attention, distraction, and cognitive control under load, Current Directions in Psychological Science, 19, 143–148, 2010.)

for us, is that it is possible to be very attentive and still miss things. One example, called *inattentional blindness*, illustrates how we can miss things even if they are clearly visible.

Inattentional Blindness

Inattentional blindness occurs when people are unaware of clearly visible stimuli if they aren't directing their attention to them (Mack & Rock, 1998). For example, Cartwright-Finch and Nilli Lavie (2007) had participants view the cross stimulus shown in **Figure 4.25**. The cross was presented for five trials, and the observer's task was to indicate which arm of the briefly flashed cross was longer, the horizontal or the vertical. This is a difficult task because the arms were just slightly different in length, the cross was flashed rapidly, and the arm that was longer changed from trial to trial. On the sixth trial, a small outline of a square was added to the display (**Figure 4.25b**). Immediately after the sixth trial, participants were asked whether they noticed if anything had appeared on the screen that they had not seen before. Out of 20 participants, only 2 reported that they had seen the square. In other words, most of the participants were "blind" to the small square, even though it was located right next to the cross.

This demonstration of inattentional blindness used a rapidly flashed geometric test stimulus. But Daniel Simons and Christopher Chabris (1999) showed that attention can affect perception within a dynamic scene by having observers view a short film that showed two "teams" of three players each. One team, dressed in white, was passing a basketball around, and the other was "guarding" that team by following them around and putting their arms up as in a basketball game (**Figure 4.26**). Observers were told to count the number of passes, a task that focused their attention on the team wearing white. After about 45 seconds, one of two events occurred: Either a woman carrying an umbrella or a person in a gorilla suit walked through the "game," an event that took 5 seconds.

After seeing the video, observers were asked whether they saw anything unusual happen or whether they saw anything other than the six players. Nearly half of the observers—46 percent—failed to report that they saw the woman or the gorilla. This experiment demonstrates that when observers are attending to one sequence of events, they can fail to notice another event, even when it is right in front of them (also see Goldstein & Fink, 1981; Neisser & Becklen, 1975).

➤ **Figure 4.26** Frame from the film shown by Simons and Chabris in which a person in a gorilla suit walks through the basketball game.

(Source: D. J. Simons & C. F. Chabris, Gorillas in our midst:Sustained inattentional blindness for dynamic events, *Perception*, 28, 1059–1074, 1999. Pion Limited, London. Figure provided by Daniel Simons.)

Inattentional Deafness

The idea that inattention can cause us to miss visual stimuli has been extended to hearing. Dana Raveh and Nilli Lavie (2015) had participants carry out a visual search task, where **visual search** involves scanning a scene to find a specific object. They presented either an easy visual search task, like the one in **Figure 4.27a**, or a hard task like the one in **Figure 4.27b**. Participants were also asked to indicate whether they heard a tone that was presented during the visual display on about a fifth of the trials. The results, shown in **Figure 4.27c**, indicate that it was more difficult to detect the tone when engaged in the hard visual search task. This situation, in which focusing on a difficult visual task results in impaired hearing, is an example of **inattentional deafness.**

This result is significant both because it shows that inattentional effects can occur across vision and hearing and also because it shows how Lavie's *load theory of attention* (see page 99) can be applied to explaining the effects of inattention. Raveh and Lavie showed that being involved in a high-load task increases the chances of missing other stimuli. Looking back at the examples of inattentional blindness in vision, we can see that the tasks involved—detecting a slight difference in line length (see Figure 4.25) or counting basketball passes (Figure 4.26)—do involve highly focused attention, so it isn't surprising that participants missed the small square or the gorilla.

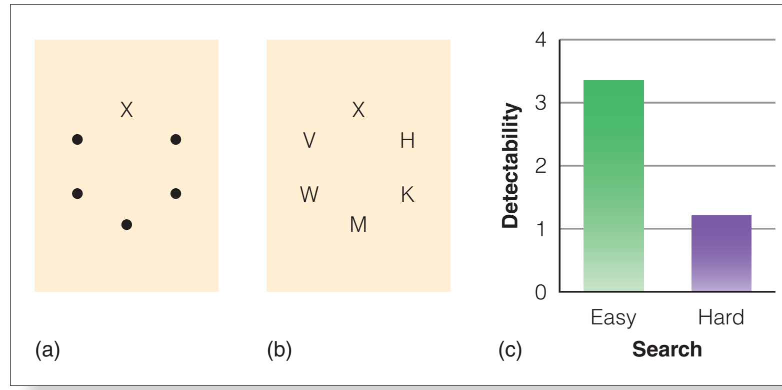

➤ **Figure 4.27** Raveh and Lavie's (2015) inattentional deafness experiment. (a) Stimuli for easy search task. Find the X. (b) Stimuli for hard search task. (c) Result, showing detectability for a tone presented during the easy and hard search tasks. The high detectability during the easy search task means the tone was easier to detect.

(Source: Based on Raveh and Lavie, 2015)

Change Detection

Researchers have also demonstrated how a lack of attention can affect perception using a procedure called **change detection**, in which one picture is presented followed by another picture, and the task is to determine what the difference is between them. To appreciate how this works, try the following demonstration before reading further.

DEMONSTRATION Change Detection

When you are finished reading these instructions, look at the picture in **Figure 4.28** for just a moment; then turn the page and see whether you can determine what is different in **Figure 4.29**. Do this now.

Were you able to see what was different in the second picture? People often have trouble detecting the change even though it is obvious when you know where to look. (Try again, paying attention to the sign near the lower-left portion of the picture.) Ronald Rensink and coworkers (1997) did a similar experiment in which they presented one picture, followed by a blank field, followed by the same picture but with an item missing, followed by a blank field. This sequence was repeated until observers were able to determine what was different about the two pictures. Rensink found that the sequence had to be repeated a number of times before the difference was detected. This difficulty in detecting changes in scenes is called **change blindness** (Rensink, 2002).

The frequency with which change blindness occurs can be startling. For example, in one study (Grimes, 1996), 100 percent of observers failed to detect a one-fourth increase in the size of a building, 92 percent failed to detect a one-third reduction in a flock of birds, 58 percent failed to detect a change in a model's swimsuit from bright pink to bright green, 50 percent failed to notice that two cowboys had exchanged their heads, and 25 percent failed to notice a 180-degree rotation of Cinderella's Castle at Disneyland!

If you find this hard to believe, you can reflect upon your own ability to detect changes while watching movies. Change blindness occurs regularly in popular films, in which some aspect of the scene, which should remain the same, changes from one shot to the next. In the *Wizard of Oz* (1939), Dorothy's (Judy Garland's)

Bruce Goldstein

➤ **Figure 4.28** Stimulus for the change detection demonstration.

➤ **Figure 4.29** Stimulus for the change detection demonstration.

hair changes length many times from short to long and back again. In *Pretty Woman* (1990), Vivian (Julia Roberts) began to reach for a croissant for breakfast that suddenly turned into a pancake. In a scene in *Harry Potter and the Sorcerer's Stone* (2001), Harry (Daniel Radcliff) suddenly changes where he is sitting during a conversation in the Great Hall. These changes in films, which are called **continuity errors**, have been well documented on the Internet (search for "continuity errors in movies").

Why does change blindness occur? The answer is that when we look at a scene in a still picture or at the ongoing action in a film, our attention is often not directed at the place where the change occurs.

What About Everyday Experience?

All of the experiments we have described—the inattentional blindness experiments, in which a distracting task kept people from noticing a test stimulus; the inattentional deafness experiment, in which focusing on a visual task results in impaired hearing; and the change blindness experiments, in which small but easily visible changes in pictures are not perceived—demonstrate that attending plays an important role in perceiving. This has implications for perception in our everyday experience, because there are a large number of stimuli present in the environment, and we are able to pay attention to only a small fraction of these stimuli at any moment. This means that we are constantly missing things in the environment.

Before you decide that our perceptual system is hopelessly flawed by its inability to detect large portions of our environment, consider the fact that we (and other animals) have somehow survived, so clearly our perceptual system is doing its job well enough to take care of most of the perceptual requirements posed by everyday life. In fact, it has been argued that the fact that our perceptual system focuses on only a small portion of the environment is one of its most adaptive features, because by focusing on what is important, our perceptual system is making optimal use of our limited processing resources.

Also even as we are focusing on what is important at the moment, our perceptual system has a warning system that responds to motion or intense stimuli, which causes us to rapidly shift our attention to things that might signal danger, such as a charging animal, a pedestrian on a collision course with us, a bright flash of light, or a loud noise. Once our attention has shifted, we can then evaluate what is happening at our new center of attention and decide whether we need to take action.

It is also important to realize that we don't need to be aware of all the details of what is happening around us. As you walk down a crowded sidewalk, you need to know where the other people are so you can avoid colliding, but you don't need to know that a particular person is wearing glasses or that another is wearing a blue shirt. You also don't need to be continually checking the details of what is happening around you because, from your past experience, you have scene schemas for city streets, country roads, or the layout of your campus that enable you to "fill in" what is around you without paying close attention (see Chapter 3, page 74).

What all of this means is that our perceptual systems are generally well adapted to take in the information we need to survive, even though we can only take in a small proportion of the information out there. But before you decide that the combination of focused attention, warning signals on the side, and filling in by schemas enables you to achieve feats of divided attention like driving and texting, remember that driving, texting, and cell phones are recent additions to the environment that weren't present when the human perceptual system evolved. Thus, as adaptive as our perceptual system might be, our modern world often puts us in situations that we are not designed to deal with and that, as we saw earlier, can lead to a dented fender, or worse.

Attention and Experiencing a Coherent World

We have seen that attention is an important determinant of what we perceive. Attention brings things to our awareness and can enhance our ability to perceive and to respond. We now consider yet another function of attention, one that is not obvious from our everyday experience. This function of attention is to help create **binding**—the process by which features such as color, form, motion, and location are combined to create our perception of a coherent object.

We can appreciate why binding is necessary by considering an everyday event: You are sitting on a bench in the park, appreciating the colors of the Fall leaves, when all of a sudden a red ball rolls across your field of view, followed closely by a small child chasing the ball. When the ball rolls by, a number of different types of cells fire in your brain. Cells sensitive to the ball's round shape fire in your temporal cortex, cells sensitive to movement fire in an area specialized for motion, and cells sensitive to depth and color fire in other areas. But even though the ball's shape, movement, depth, and color cause firing in different areas of your cortex, you don't perceive the ball as separated shape, movement, depth, and color perceptions. You experience an integrated perception of a ball, with all of the ball's features being bound together to create the coherent perception of a "rolling red ball." The question of how an object's individual features become bound together, which is called the **binding problem**, has been addressed by Anne Treisman's (1986, 1988, 1999) feature integration theory.

Feature Integration Theory

According to **feature integration theory (FIT)**, the first step in object processing is the **preattentive stage** (the first box in the flow diagram in **Figure 4.30**). As its name implies, the preattentive stage occurs *before* we focus attention on an object. Because attention is not involved, researchers argue that this stage is automatic, unconscious, and effortless. In this stage, the features of objects are analyzed independently in separate areas of the brain and are not yet associated with a specific object. For example, during the preattentive stage, the visual system of a person observing a rolling red ball would process the qualities of redness (color), roundness (form), and rightward movement (motion) separately. In the next stage of processing, called the **focused attention stage**, attention is focused on an object and the independent features are combined, causing the observer to become consciously aware of the rolling red ball.

In this two-stage process, you can think of visual features as components of a "visual alphabet." At the very beginning of the process, information about each of these components exist independently of one another, just as the letter tiles in a game of Scrabble exist as individual units when the tiles are scattered at the beginning of the game. However, just as the individual Scrabble tiles are combined to form words, the individual features combine to form perceptions of whole objects.

The idea that an object is automatically broken into features may seem counterintuitive because we always see whole objects, not objects that have been divided into individual features. The reason we aren't aware of this process of feature analysis is that it

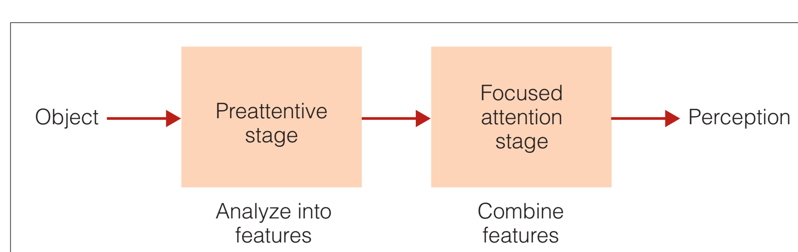

Figure 4.30 Steps in Treisman's feature integration theory. Objects are analyzed into their features in the preattentive stage, and the features are later combined with the aid of attention.

occurs early in the perceptual process, before we have become conscious of the object. Thus, when you see this book, you are conscious of its rectangular shape, but you are not aware that before you saw this rectangular shape, your perceptual system analyzed the book in terms of individual features such as lines with different orientations.

Evidence for Feature Integration Theory

To provide some perceptual evidence that objects are, in fact, analyzed into features, Anne Treisman and Hilary Schmidt (1982) did an experiment that involved a perceptual effect called *illusory conjunctions* in which one object can take on properties of another object.

Illusory Conjunctions Treisman and Schmidt presented displays like the one in **Figure 4.31**, in which four objects are flanked by two black numbers. They flashed this display onto a screen for one-fifth of a second, followed by a random-dot masking field designed to eliminate any residual perception that might remain after the stimuli were turned off. Participants were told to report the black numbers first and then to report what they saw at each of the four locations where the shapes had been. Thus, participants had to divide their a ttention across two tasks: identifying the numbers and identifying the shapes. By dividing participants' attention, Tresiman and Schmidt reduced their ability to focus attention on the shapes.

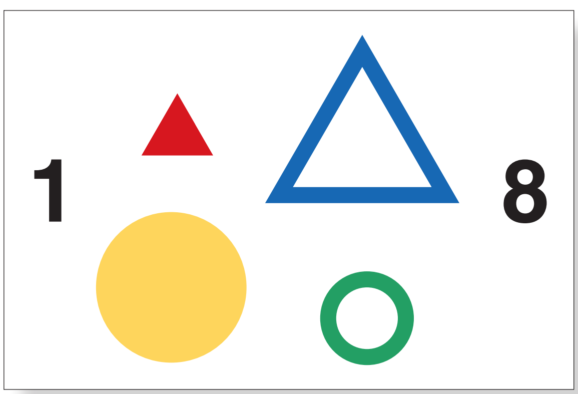

➤ **Figure 4.31** Stimuli for illusory conjunction experiment. See text for details.

(Source: A. Treisman & H. Schmidt, Illusory conjunctions in the perception of objects, *Cognitive Psychology*, 14, 107–141, 1982.)

So what did participants report seeing? Interestingly, on about one-fifth of the trials, participants reported seeing shapes that were made up of a combination of features from two different stimuli. For example, after being presented with the display in Figure 4.31, in which the small triangle is red and the small circle is green, they might report seeing a small red circle and a small green triangle. These combinations of features from different stimuli are called **illusory conjunctions**. Illusory conjunctions can occur even if the stimuli differ greatly in shape and size. For example, a small blue circle and a large green square might be seen as a large blue square and a small green circle.

Although illusory conjunctions are usually demonstrated in laboratory experiments, they can occur in other situations as well. In a class demonstration to illustrate that observers sometimes make errors in eyewitness testimony, I had a male wearing a green shirt burst into the class, grab a yellow purse that was sitting on a desk (the owner of the purse was in on the demonstration), and run out of the room. This happened so rapidly that it surprised the students in the class, whose task was to describe what had happened as eyewitnesses to a "crime." Interestingly enough, one of the students reported that a male wearing a yellow shirt grabbed a green purse from the desk! Interchanging the colors of these objects is an example of illusory conjunctions (Treisman, 2005).

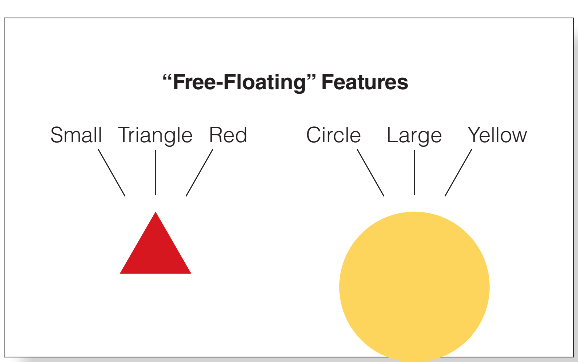

➤ **Figure 4.32** Illustration of the idea that in the preattentive stage an object's features are "free floating." Because they are not attached to a particular object, they can potentially become associated with any object in a display. When this happens, an illusory conjunction is created.

(Source: A. Treisman & H. Schmidt, Illusory conjunctions in the perception of objects, *Cognitive Psychology*, 14, 107–141, 1982.)

According to Treisman, illusory conjunctions occur because in the preattentive stage, each feature exists independently of the others. That is, features such as "redness," "curvature," or "tilted line" are, at this early stage of processing, not associated with a specific object. They are, in Treisman's (1986) words, "free floating," as shown in **Figure 4.32**, and can therefore be incorrectly combined if there is more than one object, especially in laboratory situations when briefly flashed stimuli are followed by a masking field.

When I describe this process in class, some students aren't convinced. One student said, "I think that when people look at an object, they don't break it into parts. They just see what they see." To convince such students (and the many others who, at the beginning of the course, are not comfortable with the idea that perception sometimes involves rapid processes we aren't aware of), I describe the case of R.M., a patient who had parietal lobe damage that resulted in a condition called **Balint's syndrome**. A crucial characteristic of Balint's syndrome is an inability to focus attention on individual objects.

According to feature integration theory, lack of focused attention would make it difficult for R.M. to combine features correctly, and this is exactly what happened. When R.M. was presented with two different letters of different colors, such as a red T and a blue O, he reported illusory conjunctions such as "blue T" on 23 percent of the trials, even when he was able to view the letters for as long as 10 seconds (Friedman-Hill et al., 1995; Robertson et al., 1997). The case of R.M. illustrates how a breakdown in the brain can reveal processes that are not obvious when the brain is functioning normally.

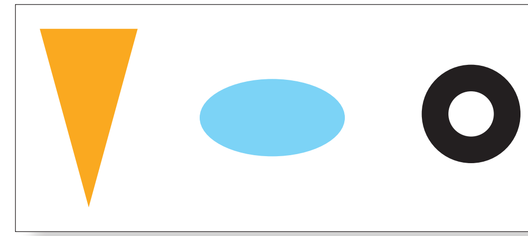

➤ **Figure 4.33** Stimuli used to show that top-down processing can reduce illusory conjunctions. (Source: A. Treisman & H. Schmidt, Illusory conjunctions in the perception of objects, *Cognitive Psychology*, 14, 107–141, 1982.)

The feature analysis approach involves mostly bottom-up processing because knowledge is usually not involved. In some situations, however, top-down processing can come into play. For example, when Treisman and Schmidt (1982) did an illusory conjunction experiment using stimuli such as the ones in **Figure 4.33** and asked participants to identify the objects, the usual illusory conjunctions occurred; the orange triangle, for example, would sometimes be perceived to be black. However, when she told participants that they were being shown a carrot, a lake, and a tire, illusory conjunctions were less likely to occur, and participants were more likely to perceive the triangular "carrot" as being orange. In this situation, the participants' knowledge of the usual colors of objects influenced their ability to correctly combine the features of each object. In our everyday experience, in which we often perceive familiar objects, top-down processing combines with feature analysis to help us perceive things accurately.

Visual Search Another approach to studying the role of attention in binding has used a type of visual search task called a **conjunction search.**

DEMONSTRATION Searching for Conjunctions

We can understand what a conjunction search is by first describing another type of search called a feature search. Before reading further, find the horizontal line in **Figure 4.34a**. This is a feature search because you could find the target by looking for a single feature—"horizontal." Now find the green horizontal line in **Figure 4.34b**. This is a conjunction search because you had to search for a combination (or conjunction) of two or more features in the same stimulus—"horizontal" and "green." In Figure 4.34b, you couldn't focus just on green because there are vertical green lines, and you couldn't focus just on horizontal because there are horizontal red lines. You had to look for the conjunction of horizontal *and* green.

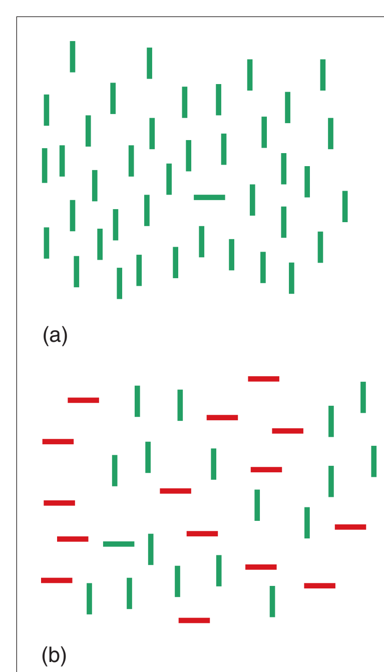

➤ **Figure 4.34** Find the horizontal line in (a) and then the green horizontal line in (b). Which task took longer?

Conjunction searches are useful for studying binding because finding the target in a conjunction search involves scanning a display in order to focus attention at a specific location. To test the idea that attention to a location is required for a conjunction search, a number of researchers have tested R.M., the Balint's patient, and have found that he cannot find the target when a conjunction search is required (Robertson et al., 1997). This is what we would expect because of R.M's difficulty in focusing attention. R.M. can, however, find targets when only a feature search is required, as in Figure 4.34a, because attention-at-a-location is not required for this kind of search. Visual scanning experiments, both on R.M. and normal observers, provides evidence that supports the idea that attention is an essential component of the mechanism that creates our perception of objects from a number of different features (Wolfe, 2012).

SOMETHING TO CONSIDER

Attentional Networks

We've seen how attention can affect responding in the brain by enhancing activity at a location (page 107) or by expanding the area dedicated to a specific type of objects (page 108).

But we need to take a step further to fully understand the connection between attention and the brain. We need to look at how the brain is set up to make attention work. To do that, we consider the neural networks that transmit signals across many areas of the brain, which we introduced in Chapter 2 (see page 45).

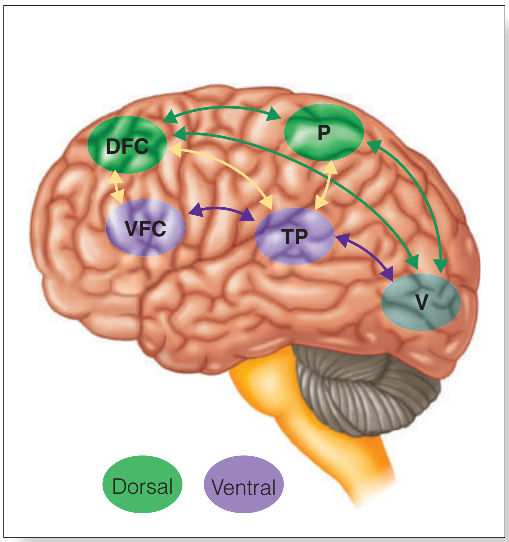

Figure 4.35 Main structures in the two attention networks. V = visual cortex. Dorsal attention network: P = parietal cortex; DFC = dorsal frontal cortex; Ventral attention network: TP = junction of temporal and parietal lobes; VFC = ventral frontal cortex.

(Source: Based on Vossel et al., 2014, Figure 1)

Neuroimaging research has revealed that there are neural networks for attention associated with different functions. Consider, for example, how attention is directed by scanning a scene (page 103). We saw that attention is determined by stimulus salience—physical properties of stimuli—and by higher-level top-down functions such as scene schemas, as when an unusual object appears in a scene, or task demands, as in the example of making a peanut butter and jelly sandwich. Imaging experiments in which participants carried out tasks involving salience or involving top-down processes have revealed two different networks: the **ventral attention network**, which controls attention based on salience, and the **dorsal attention network**, which controls attention based on top-down processes (**Figure 4.35**).

Identifying different networks for different functions was a big step toward understanding how the brain controls attention. But researchers have gone beyond just identifying networks to looking at the dynamics of how information flows in these networks. Remember the helicopter ride we took in Chapter 2 to observe the flow of traffic in a network of city streets, which represented a neural network (page 49). We noted that traffic flow changed depending on changing conditions. For example, traffic flow toward the stadium increased on the weekend of the big football game. Similarly, flow in attention systems changes depending on whether attention is being controlled by stimulus salience or by top-down factors, with more flow in the ventral network for control by salience and more in the dorsal network when flow is controlled by top-down factors.

But to fully understand the dynamic nature of attention, we need to go one step further. Different tasks don't just shift activity from one pathway to another. They also change the **effective connectivity** between different areas in a network. Effective connectivity refers to how easily activity can travel along a particular pathway.

We can illustrate effective connectivity by returning to our road network example, in which we noted that the flow of traffic is directed more toward the stadium on the day of the football game. Sometimes, when situations like this occur, the people in charge of regulating traffic in the city open up more lanes moving toward the stadium before the game, and then open more lanes moving away from the stadium after the game. In other words, the basic road system remains the same, but the flow becomes easier in certain directions, depending on conditions.

This is exactly what happens for attention when the effective connectivity between different structures in a network changes depending on conditions. How does this effective connectivity change? One mechanism that has been suggested, **synchronization**, is illustrated by the results of an experiment by Conrado Bosman and coworkers (2012) in which they recorded a response called the *local field potential (LFP)* from a monkey's cortex. LFPs, which are recorded by small disc electrodes placed on the surface of the brain, record signals from thousands of neurons near the electrode. LFP responses were recorded from an electrode at A on the brain, where signals from the visual stimulus arrive. Signals are also recorded from an electrode at B on the brain, which is connected to A and so receives signals from A (**Figure 4.36a**).

Bosman found that the visual stimulus caused an LFP response at A in the cortex and also at B, because A sends signals to B. He also found that when the monkey wasn't paying attention to the visual stimulus, the responses recorded from A and B were unsynchronized (**Figure 4.36b**). However, when the monkey focused attention on the visual stimulus, the signals from A and B become synchronized (**Figure 4.36c**). It has been hypothesized that synchronization such as this results in more effective communication between the two areas (see Bosman et al., 2012; Buschman & Kastner, 2015).

In addition to the ventral and dorsal attention networks, another network, called the **executive attention network**, has been proposed. This network is extremely complex and may involve two separate networks (Petersen & Posner, 2012). Rather than list all the structures involved, let's focus on what the executive attention network does.

The executive attention network is responsible for executive functions. **Executive functions** include a range of processes that involve controlling attention and dealing with conflicting responses. One example is the Stroop test (see page 100), in which the task involves focusing on the color of the ink and ignoring the color spelled out by the words. But executive attention extends into real life as well, any time there is conflict between different possible courses of action.

Dealing with conflict in everyday life has been called a number of things, including **cognitive control**, **inhibitory control**, and **willpower**. You can probably think of situations in which you were faced with a temptation that was difficult to resist. If so, your executive attention system was involved in dealing with this situation. As stated in the title of the song "Should I Stay or Should I Go," by the British punk rock group The Clash, decisions and temptation are a part of life. In the next chapter, we will see that there are connections between attention, cognitive control, and a type of memory called working memory.

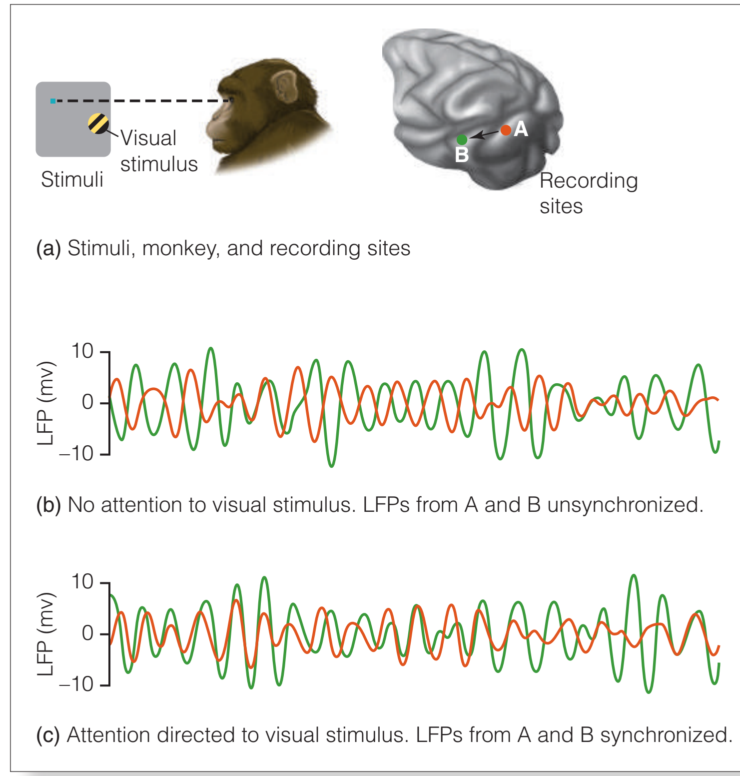

➤ **Figure 4.36** Bosman et al. (2012) demonstration of synchronization caused by attention. (a) The monkey is looking at the blue dot on the screen. Local field potentials are being recorded from locations A and B on the cortex, which are connected. (b) When the monkey is not paying attention to the visual stimulus, the LFP responses from A and B are unsynchronized. (c) When the monkey focuses its attention on the visual stimulus, the LFP responses of A and B become synchronized.

(Source: Figure courtesy of Pascal Fries and Conrado Bosman)

•••

The story we have told in the last two chapters has been about interacting with things in the environment. We perceive objects visually, hear sounds, experience smells or someone touching us, and in some cases we pay attention to some of these things more than others. Both perception and attention support our ability to know about our environment and to act within it. But to take us beyond having immediate experiences, we need to be able to store some of what is happening to us so we can remember it later. This function is achieved by the process of memory, which not only helps us survive but also determines our identity as a person. This is so important that we will spend the next four chapters discussing the process of memory. As you will see, many of the things we have introduced in our discussion of perception and attention—the principle of representation, the importance of knowledge gained from experience, the way we use inference and prediction, and our active interaction with both ideas and things—are central to our understanding of memory.

TEST YOURSELF 4.3

1. Describe Schneider and Shiffrin's experiment that demonstrated automatic processing. What are some real-life examples of automatic processing? When is automatic processing not possible?
2. What conclusions can be reached from the results of experiments testing the ability to drive while talking on a cell phone?
3. What is the evidence that cell phones can affect performance in situations in addition to driving?
4. How can a principle of operant conditioning explain why some people check their cell phones so often?
5. What is mind wandering, and how does it affect the ability to focus attention on tasks? What brain network is associated with mind wandering?
6. Describe the following evidence that attention is sometimes necessary for perception: the inattentional blindness experiment; the "basketball-passing" experiment; the change detection experiments.
7. What is inattentional deafness, and what does the inattentional deafness experiment described in the text tell us about the relation between load theory and the effects of inattention?
8. Why can we say that we don't need to be aware of all of the details of what is happening around us?
9. What is binding, and why is it necessary? What is the binding problem?
10. Describe Treisman's feature integration theory. What does the theory seek to explain about perceiving objects? What are the stages of the theory, and at what point does attention become involved?
11. What are illusory conjunctions, and what do they demonstrate about feature analysis? How have illusory conjunction experiments supported the role of attention in feature analysis? How do experiments with Balint's syndrome patients support feature integration theory?
12. What is a feature search? A conjunction search? Which type of search did the Balint's patient find difficult? What does that tell us about the role of attention in feature integration?
13. Describe how attention is controlled by different types of attentional networks. Be sure you understand the functions of the dorsal attention network, the ventral attention network, and the executive attention network, and the principles of effective connectivity and synchronization.

CHAPTER SUMMARY

1. Selective attention, the ability to focus on one message while ignoring all others, has been demonstrated using the dichotic listening procedure.
2. A number of models have been proposed to explain the process of selective attention. Broadbent's filter model proposes that the attended message is separated from the incoming signal early in the analysis of the signal. Treisman's model proposes later separation and adds a dictionary unit to explain how the unattended message can sometimes get through. Late selection models propose that selection doesn't occur until messages are processed enough to determine their meaning.
3. Lavie proposes that our ability to ignore distracting stimuli can be explained by considering processing capacity and perceptual load. Her load theory of attention states that distraction is less likely for high-load tasks because no capacity remains to process potential distracting stimuli.
4. The Stroop effect demonstrates how a powerful task-irrelevant stimulus, such as meaningful words that result in a response that competes with the observer's task, can capture attention.