

NINTH EDITION

SENSATION

and

PERCEPTION

E. Bruce Goldstein



Chris Trotman/PCN/Corbis

Taking Action

CHAPTER CONTENTS

The Ecological Approach to Perception

The Moving Observer Creates Information in the Environment

Self-Produced Information

The Senses Do Not Work in Isolation

Navigating Through the Environment

Do Observers Use Optic Flow Information?

Driving a Car

Walking

Wayfinding

Acting on Objects

Affordances: What Objects Are Used For

The Physiology of Reaching and Grasping

Observing Other People's Actions

Mirroring Others' Actions in the Brain

Predicting People's Intentions

SOMETHING TO CONSIDER: Action-Based Accounts of Perception

Think About It

VL The Virtual Lab icons direct you to specific animations and videos designed to help you visualize what you are reading about. Virtual Labs are listed at the end of the chapter, keyed to the page on which they appear, and can be accessed through Psychology CourseMate.

◀ How did McKayla Maroney of the U.S. gymnastics team, vaulting at the 2012 London Olympics, get into this position, and how did she execute a successful landing just moments later? As we will see in this chapter, the answer involves a close connection between perception and action, and this connection holds not just for spectacular athletic feats, but also for everyday actions such as walking across campus or reaching across a table to pick up a cup of coffee.

Some Questions We Will Consider:

- What is the connection between perceiving and moving through the environment? (p. 154)
- What is the connection between somersaulting and vision? (p. 155)
- How do neurons called mirror neurons respond when a person perceives an action and when the person watches someone else perceive the same action? (p. 166)

Serena straps on her helmet for what she anticipates will be a fast, thrilling, and perhaps dangerous ride. As an employee of the Speedy Delivery Package Service, her mission is to deliver the two packages strapped to the back of her bicycle to an address 30 blocks uptown. Once on her bike, she weaves through traffic, staying alert to close calls with cars, trucks, pedestrians, and potholes. Seeing a break in traffic, she reaches down to grab her water bottle to take a quick drink before having to deal with the next obstacle. “Yes,” Serena thinks, “I can multitask!” As she replaces the water bottle, she downshifts and keeps a wary eye out for the pedestrian ahead who looks as though he might decide to step off the curb at any moment.

Serena faces a number of challenges that involve both perception—using her sight and hearing to monitor what is happening in her environment—and action—staying balanced on her bike, staying on course, reaching for her water bottle, and being ready to avoid the pedestrian who does, as Serena predicted, step off the curb just as she is approaching.

We have discussed some of these things in the last two chapters: perceiving a scene and individual objects within it, scanning the scene to shift attention from one place to another, focusing on what is important and ignoring what is not, and relying on prior knowledge about characteristics of the environment. This chapter takes all of these things a step further by considering the processes involved in being *physically active* and interacting with objects within a scene. In other words, we are taking perception out into the world, where perception often occurs “on the run,” as in Serena’s bike trip, or in a more

relaxed setting, as when Serena, resting in a coffee shop after her ride, reaches across the table to pick up her coffee cup. As we explain how Serena is able to stay on course, grab her water bottle, predict what is going to happen ahead, and reach for her cup of coffee, we will be describing how perceiving and taking action interact with one another. We will see, in this chapter, that we need to consider action to truly understand perception. To begin our discussion of perception and action, we consider an early and influential approach proposed by J. J. Gibson, who founded the ecological approach to perception.

The Ecological Approach to Perception

During World War II, J. J. Gibson studied the kind of perceptual information that airplane pilots use when coming in for a landing. In his first book, *The Perception of the Visual World* (1950), Gibson proposed that pilots use information that is created by their own movement. What this means is that they look out the window and, because of their movement, the terrain is rushing by beneath them. The perceived movement of the terrain provides information that helps the pilot guide the plane in for a landing. We will consider how pilots might use this information in a moment, but first it is important to note the difference between Gibson's approach and the way perception was being studied in the mid-20th century (Goldstein, 1981).


From the 1950s until the 1980s, the dominant way perception research was carried out was by having stationary observers look at stimuli in a laboratory situation. Gibson's idea was that this traditional way of studying perception couldn't explain perception as experienced by moving observers, such as pilots landing an airplane or people riding a bike or walking down the street. The correct approach, suggested Gibson, was to study how people perceive as they move through the environment. This focus on observers moving through the environment was the starting point for the **ecological approach to perception**. The ecological approach focuses on studying moving observers and on determining how their movement creates perceptual information that both guides further movement and helps observers perceive the environment.

The Moving Observer Creates Information in the Environment

To understand what it means to say that movement creates perceptual information, imagine that you are driving down the street. No other cars or people are visible, so everything around you—buildings, trees, traffic signals—is stationary. But even though the objects are stationary, your movement *relative to the objects* causes you to see the houses and trees moving past when you look out the side window. And when you look at the road ahead, you see the road moving toward the front of your car. As your car hurtles forward when crossing a bridge, everything around you—the sides and top of the bridge and

the road below—moves past you in a direction opposite to the direction you are moving (**Figure 7.1**). All of the movement you are seeing is called **optic flow**. According to Gibson, optic flow provides information about how rapidly we are moving and where we are headed. Optic flow has two characteristics:

1. Optic flow is more rapid near the moving observer, as indicated by the length of the arrows in Figure 7.1, with longer arrows indicating more rapid flow. The different speed of flow—fast near the observer and slower farther away—is called the **gradient of flow**. According to Gibson, the gradient of flow provides information about how fast the observer is moving.
2. There is no flow at the destination toward which the observer is moving. The absence of flow at the destination point is called the **focus of expansion (FOE)**. In Figure 7.1 the FOE, marked by the dot, is at the end of the bridge, and in **Figure 7.2**, which shows optic flow lines for an airplane coming in for a landing, the FOE is indicated by a small red dot. The FOE indicates the place where the plane will touch down on the runway if it maintains its present course. **VL**



Bruce Goldstein

Figure 7.1 The side and top of the bridge and the road below appear to move toward a car that is moving forward. This movement is called optic flow.




Figure 7.2 Optic flow created by an airplane coming in for a landing. The focus of expansion (FOE), indicated by the red dot, is the place where the plane will touch down on the runway. From Gibson, J. J.

The perception of the visual world. Boston: Houghton Mifflin, 1950. Figure 58, page 128.

Another important concept of the ecological approach is the idea of **invariant information**—information that remains constant even when the observer is moving. Optic flow provides invariant information because flow information is present as long as the observer is moving through the environment. Of course, as the observer moves through a scene, the flow might look different—houses flow past on a city street, and trees on a country road—but flow is still there.

The FOE is another invariant property because it always occurs at the point toward which the observer is moving. If an observer changes direction, the FOE shifts to a new location, but the FOE is still there. Thus, even when specific aspects of a scene change, flow and the FOE continue to provide information about how fast a person is moving and where he or she is heading. When we consider depth perception in Chapter 11, we will see that Gibson proposed other sources of invariant information that indicate an object’s size and its distance from the observer.

Self-Produced Information

Another idea of the ecological approach is **self-produced information**: When a person makes a movement, that movement creates information, and this information is, in turn, used to guide further movement (**Figure 7.3**). For example, when a person is driving down the street, movement of the car provides flow information, and the observer then uses this flow information to help steer the car in the right direction. Another example of movement that creates information that is used to guide further movement is provided by somersaulting.

We can appreciate the problem facing a gymnast who wants to execute an airborne backward somersault (or back flip) by realizing that, within 600 ms, the gymnast must execute the somersault and then end in exactly the correct body configuration precisely at the moment that he or she hits the ground (**Figure 7.4**). One way this could be accomplished is to learn to run a predetermined sequence of motions within a specific period of time. In this case, performance should be the same with eyes open or closed. However, Benoit Bardy and Makel Laurent (1998) found that expert gymnasts performed somersaults better with their eyes open. Films showed that

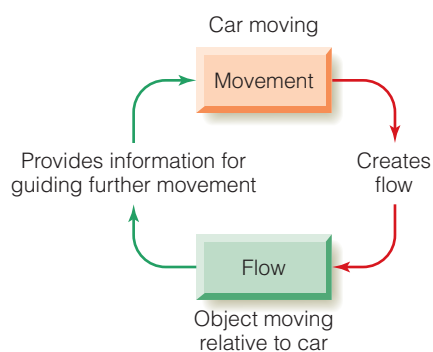


Figure 7.3 The relationship between movement and flow is reciprocal, with movement causing flow and flow guiding movement. This is the basic principle behind much of our interaction with the environment. © Cengage Learning 2014



Figure 7.4 “Snapshots” of a somersault, or backflip, starting on the left and finishing on the right. From Bardy, B. G., & Laurent, M. (1998). How is body orientation controlled during somersaulting? *Journal of Experimental Psychology: Human Perception and Performance*, 24, 963–977. Copyright © 1998 by The American Psychological Society. Reprinted by permission.

when their eyes were open, the gymnasts appeared to be making in-the-air corrections to their trajectory. For example, a gymnast who initiated the extension of his or her body a little too late compensated by performing the rest of the movement more rapidly.

Another interesting result was that closing the eyes did not affect the performance of novice somersaulters as much as it affected the performance of experts. Apparently, experts learn to coordinate their movements with their perceptions, but novices have not yet learned to do this. Therefore, when the novices closed their eyes, the loss of visual information had less effect than it did for the experts. Thus, somersaulting, like driving a car or piloting an airplane, involves using information created by movement to guide further movement.

The Senses Do Not Work in Isolation

Gibson also proposed that the senses do not work in isolation. He believed that rather than considering vision, hearing, touch, smell, and taste as separated senses, we should consider how each one provides information for the same behaviors. One example of how a behavior originally thought to be the exclusive responsibility of one sense is also served by another one is provided by the sense of balance.

Your ability to stand up straight and to keep your balance while standing still or walking depends on systems that enable you to sense the position of your body. These systems include the vestibular canals of your inner ear and receptors in the joints and muscles. However, Gibson argued that information provided by vision also plays a role in keeping our balance. One way to illustrate the role of vision in balance is to consider what happens when visual information isn’t available, as in the following demonstration.

DEMONSTRATION Keeping Your Balance

Keeping your balance is something you probably take for granted. Stand up. Raise one foot from the ground and stay balanced on the other. Then close your eyes and notice what happens.

Did staying balanced become more difficult when you closed your eyes? Vision provides a frame of reference that helps the muscles constantly make adjustments to help maintain balance.

The importance of vision in maintaining balance was demonstrated by David Lee and Eric Aronson (1974). Lee and Aronson placed 13- to 16-month-old toddlers in a “swinging room” (Figure 7.5). In this room, the floor was stationary, but the walls and ceiling could swing toward and away from the toddler. Figure 7.5a shows the room swaying toward the toddler. This movement of the wall creates the optic flow pattern on the right. Notice that this pattern is similar to the optic flow that occurs when moving forward, as when you are driving through a tunnel.

The flow pattern that the toddler observes creates the impression that he or she is swaying forward. This causes the toddler to sway back to compensate (Figure 7.5b). When the room moves back, as in Figure 7.5c, the flow pattern creates the impression of swaying backward, so the toddler sways forward to compensate. Although a few of the toddlers

were unaffected by the sway, 26 percent swayed, 23 percent staggered, and 33 percent fell down, even though the floor remained stationary throughout the entire experiment!

Adults were also affected by the swinging room. Some of them braced themselves so they just swayed back and forth rather than staggering or falling down. Lee describes their behavior as follows: “oscillating the experimental room through as little as 6 mm caused adult subjects to sway approximately in phase with this movement. The subjects were like puppets visually hooked to their surroundings and were unaware of the real cause of their disturbance” (p. 173). Adults who didn’t brace themselves could, like the toddlers, be knocked over by their perception of the moving room.

The swinging room experiments show that vision is such a powerful determinant of balance that it can override the traditional sources of balance information provided by the inner ear and the receptors in the muscles and joints (see also Fox, 1990). In a developmental study, Bennett Berthenthal and coworkers (1997) showed that infants as young as 4 months old sway back and forth in response to movements of a room,




Figure 7.5 Lee and Aronson’s swinging room. (a) Moving the wall toward the observer creates an optic flow pattern associated with moving forward, so (b) the observer sways backward to compensate. (c) As the wall moves away from the observer, flow corresponds to moving backward, so the person leans forward to compensate and may even lose his or her balance. Based on Lee, D. N., & Aronson, E. (1974). Visual proprioceptive control of standing in human infants. *Perception and Psychophysics*, 15, 529–532, Figure 2.

and that the coupling of the room's movement and the swaying becomes closer with age. (See also Stoffregen et al., 1999, for more evidence that flow information can influence posture while standing still; and Warren et al., 1996, for evidence that flow is involved in maintaining posture while walking.)

Gibson's emphasis on (1) the moving observer, (2) identifying invariant information in the environment that observers use for perception, and (3) considering the senses as working together was revolutionary for its time. But even though perception researchers were aware of Gibson's ideas, most research continued in the traditional way—testing stationary subjects looking at stimuli in laboratory settings. Of course, there is nothing wrong with testing stationary observers in the laboratory, and much of the research described in this book takes this approach. However, Gibson's idea that perception should also be studied as it is often experienced (by observers who are moving and in more naturalistic settings) finally began to take hold in the 1980s, and today perception in naturalistic settings is one of the major themes of perception research.

In the remainder of this chapter we will consider the following ways that perception and action occur together in the environment: (1) navigating through the environment by walking or driving; (2) interacting with objects in the environment by reaching out and grasping them; and (3) watching other people take action in the environment.

Navigating Through the Environment

Gibson proposed that optic flow provides information about where a moving observer is heading. But can observers actually use this information? We consider this question next and then consider sources of information in addition to optic flow that help people navigate through the environment.

Do Observers Use Optic Flow Information?

Research on whether people use flow information has asked observers to make judgments regarding where they are heading based on computer-generated displays of moving dots that create optic flow stimuli. The observer's task is to judge, based on optic flow stimuli, where he or she would be heading relative to a reference point such as the vertical line in **Figures 7.6a** and **b**. The flow in Figure 7.6a indicates movement directly toward the line, and the flow in Figure 7.6b indicates movement to the right of the line. Observers viewing stimuli such as this can judge where they are heading relative to the vertical line to within about 0.5 to 1 degree (Warren, 1995, 2004; also see Fortenbaugh et al., 2006; Li, 2006). **VL**

Psychophysical results such as these support Gibson's idea that optic flow provides information about where a person is heading. Researchers have also identified neurons




Figure 7.6 (a) Optic flow generated by a person moving straight ahead toward the vertical line on the horizon. The lengths of the lines indicate the person's speed. (b) Optic flow generated by a person moving in a curved path that is headed to the right of the vertical line. From Warren, W. H. (1995). *Self-motion: Visual perception and visual control*. In W. Epstein & S. Rogers (Eds.), *Handbook of perception and cognition: Perception of space and motion* (pp. 263–323). Copyright © 1965, with permission from Elsevier.

in the brain that respond to flow patterns. One place where these neurons are found is in the medial superior temporal area (MST), which we will see in Chapter 8 is important for perceiving movement (**Figure 7.7**).




Figure 7.7 The human brain, showing the medial superior temporal area (MST), which responds to optic flow, as discussed here. Other areas, which will be discussed later, are the parietal reach region (PRR) in the parietal lobe, which is involved in reaching and grasping, and the premotor cortex (PM), which is involved in observing other people's actions. © Cengage Learning 2014

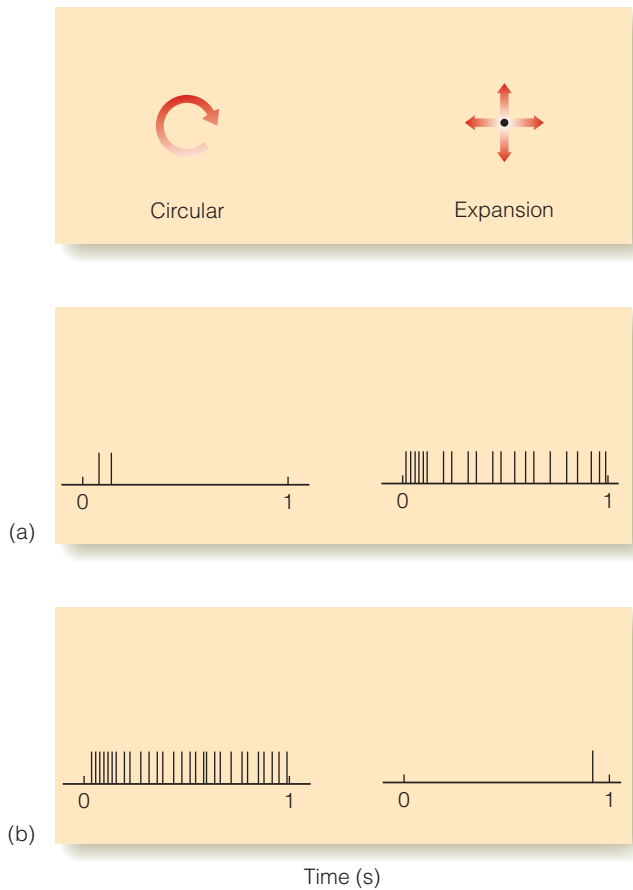



Figure 7.8 (a) Response of a neuron in the monkey's MST that responds to an expanding stimulus, but hardly responds to a stimulus that moves in a circular motion. (b) A neuron that responds to circular movement, but doesn't respond to expansion. Based on Graziano, M. S. A., Andersen, R. A., & Snowden, R. J. (1994). Tuning of MST neurons to spiral motions. *Journal of Neuroscience*, 14, 54-67.


Figure 7.8 shows the response of a neuron in a monkey's MST that responds best as the monkey observes a pattern of dots that are expanding outward, as would occur if the monkey were moving forward (**Figure 7.8a**), and another neuron that responds best to circular motions, as would occur if the monkey were swinging through the trees (**Figure 7.8b**; see also Duffy & Wurtz, 1991; Orban et al., 1992; Raffi et al., 2002; Regan & Cynader, 1979). What does the existence of these optic flow neurons mean? We know from previous discussions that finding a neuron that responds to a specific stimulus is only the first step in determining whether this neuron has anything to do with perceiving that stimulus (see Chapter 3, page 66). The next step is to demonstrate a connection between the neuron's response and behavior.

Kenneth Britten and Richard van Wezel (2002) demonstrated a connection between the response of neurons in MST and behavior by first training monkeys to indicate whether the flow of dots on a computer screen indicated movement to the left or right of straight ahead. For example, **Figure 7.9** shows a monkey viewing a flow that would occur if the monkey were moving slightly to the left.

The left bar in **Figure 7.9b** shows that the monkey responded to a stimulus like this by judging the movement as



(a)



(b)

Figure 7.9 (a) A monkey watches a display of moving dots on a computer monitor. The dots indicate the flow pattern for movement slightly to the left of straight ahead. (b) Effect of microstimulation of the monkey's MST neurons that were tuned to respond to leftward movement. Stimulation (red bar) increases the monkey's judgment of leftward movement. Based on data from Britten, K. H., & van Wezel, R. J. A. (2002). Area MST and heading perception in macaque monkeys. *Cerebral Cortex*, 12, 692-701.

being to the left on 60 percent of the trials. But if, as the monkey was making its judgment, Britten and van Wezel electrically stimulated MST neurons that were tuned to respond to flow associated with movement to the left, the monkey's judgment was shifted even more to the left, increasing from 60 percent to 80 percent of the trials. This demonstration that stimulating flow neurons can influence the monkey's judgment of the direction of movement supports the idea that flow neurons can, in fact, help determine the direction of perceived movement.

Driving a Car

The experiments described above show that observers and neurons can respond to the flow indicated by computer-generated patterns of moving dots. But what about the flow that occurs in an actual environmental situation such as driving? To study information people use to stay on course when driving, Michael Land and David Lee (1994) fitted an automobile with instruments to record the angle of the




Figure 7.10 Results of Land and Lee's (1994) experiment. The ellipses indicate the place where the drivers were most likely to look while driving down (a) a straight road and (b) a curve to the left. From Land, M. F., & Lee, D. N. (1994). Where we look when we steer. *Nature*, 377, 742–744.

steering wheel and the speed, and measured where the driver was looking with a video eye tracker. According to Gibson, the focus of expansion (FOE) provides information about the place toward which a moving observer is headed. However, Land and Lee found that although drivers look straight ahead while driving, they tend to look at a spot in front of the car rather than looking directly at the FOE (**Figure 7.10a**).

Land and Lee also studied where drivers look as they are negotiating a curve. This task poses a problem for the idea of FOE because the driver's destination keeps changing as the car rounds the curve. Land and Lee found that when going around a curve, drivers don't look directly at the road, but instead look at the tangent point of the curve on the side of the road, as shown in **Figure 7.10b**. Because drivers don't look at the FOE, which would be in the road directly ahead, Land and Lee suggested that drivers probably use information in addition to optic flow to determine the direction they are heading. An example of this additional information would be noting the position of the car relative to the lines in the center of the road or relative to the side of the road. (See also Kandil et al., 2009; Land & Horwood, 1995; Rushton & Salvucci, 2001; Wann & Land, 2000; Wilkie & Wann, 2003, for more research on the information drivers use to stay on the road.) **VL**

Walking

How do people navigate on foot? Apparently, an important strategy used by walkers (and perhaps drivers as well) that does not involve optic flow is the **visual direction strategy**, in which observers keep their body pointed toward a target. If they go off course, the target will drift to the left or right (**Figure 7.11**). When this happens, the walker can correct course by recentering the target (Fajen & Warren, 2003; Rushton et al., 1998).

Another indication that flow information is not always necessary for navigation is that we can find our way even when flow information is minimal, such as at night or in a snowstorm (Harris & Rogers, 1999). Jack Loomis and coworkers (Loomis et al., 1992; Philbeck, Loomis, & Beall, 1997) have demonstrated this by eliminating flow altogether, with a “blind walking” procedure in which people observe a target object located up to 12 meters away, then walk to the target with their eyes closed. **VL**

These experiments show that people are able to walk directly toward the target and stop within a fraction of a meter of it (red lines in **Figure 7.12**). In fact, people can do this even when they are asked to walk off to the side first and then make a turn and walk to the target, while keeping their eyes closed. Some records from these “angled” walks are shown by the blue lines in Figure 7.12, which depict the paths taken

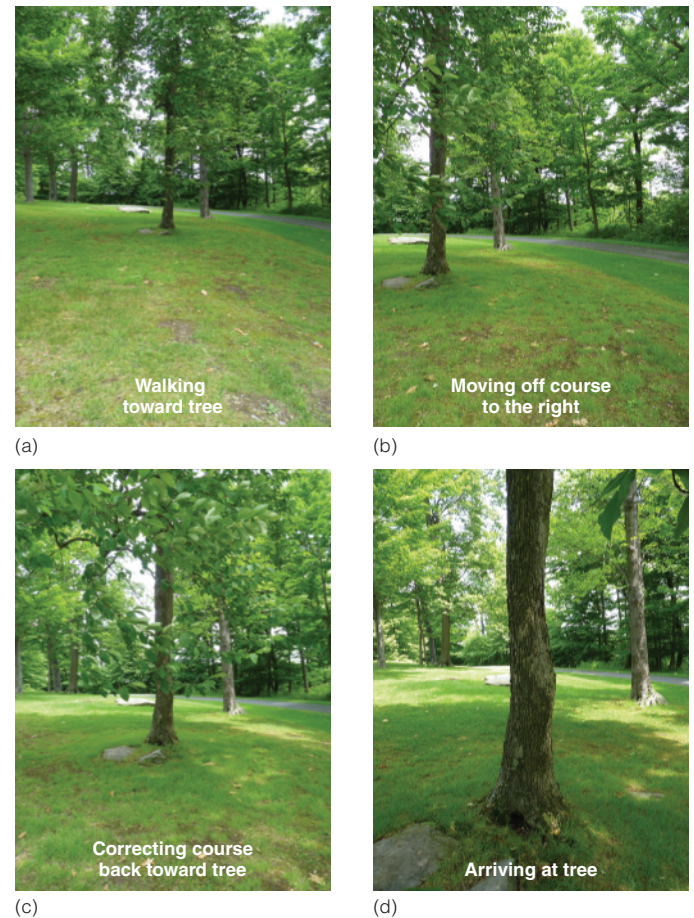


Figure 7.11 (a) As long as a person is moving toward the tree, it remains in the center of the person's field of view. (b) When the person walks off course, the tree drifts to the side. (c) When the person corrects course, the tree moves back to the center of the field of view, until (d) the person arrives at the tree.




Figure 7.12 The results of a “blind walking” experiment (Philbeck et al., 1997). Participants looked at the target, which was 6 meters from the starting point, then closed their eyes and begin walking to the left. They turned either at point 1 or 2, keeping their eyes closed the whole time, and continued walking until they thought they had reached the target. © Cengage Learning 2014

when a person first walked to the left from the “start” position and then was told to turn either at turning point 1 or 2 and walk to a target that was 6 meters away. The fact that the person generally stopped close to the target shows that we are able to accurately navigate short distances in the absence of any visual stimulation at all (also see Sun et al., 2004).

Wayfinding

So far we have been considering information that observers might use to travel toward a destination they can see. But we often travel to destinations we can’t see from the starting point, such as when we walk across campus from one class to another or drive to a destination several miles away. This kind of navigation, in which we take a route that involves making turns, is called **wayfinding**.

Our ability to get from one place to another may seem simple, especially for routes you have traveled many times. But just as there is nothing simple about perception, there is nothing simple about wayfinding. It is a complex process that involves perceiving objects in the environment, remembering objects and their relation to the overall scene, and knowing when to turn and in what direction.

The Importance of Landmarks One important source of information for wayfinding is **landmarks**—objects on the route that serve as cues to indicate where to turn. Sahar Hamid and coworkers (2010) studied how subjects used landmarks as they learned to navigate through a mazelike environment displayed on a computer screen in which pictures of common objects served as landmarks. Subjects first navigated through the maze until they learned its layout (training phase) and then were told to travel from one location in the maze to another (testing phase). During both the training and testing phases, subjects’ eye movements were measured using a head-mounted eye tracker like the one used in the

experiment described in Chapter 6 in which eye movements were measured as a subject made a peanut butter and jelly sandwich (see page 132). This maze contained both *decision-point landmarks*—objects at corners where the subject had to decide which direction to turn—and *non-decision-point landmarks*—objects located in the middle of corridors that provided no information about how to navigate.

The eye-tracking measurements showed that subjects spent more time looking at decision-point landmarks than at non-decision-point landmarks, probably because the decision-point landmarks were more important for navigating the maze. In fact, when maze performance was tested with half of the landmarks removed, removing landmarks that had been viewed less (and were likely to be in the middle of the corridors) had little effect on performance (**Figure 7.13a**). However, removing landmarks that observers had looked at longer caused a substantial drop in performance (**Figure 7.13b**).

It makes sense that landmarks that are looked at the most would be the ones that are used to guide navigation. Another study, in which subjects learned a walking route through the University of Pennsylvania campus, showed that after subjects had learned the route, they were more likely to recognize pictures of buildings that were located at decision points than those located in the middle of the block (Schinazi & Epstein, 2010).

The studies we have described have measured eye movements, maze performance, and recognition, all of which are behaviors related to landmarks. But what is happening in the brain? When subjects in the University of Pennsylvania study were shown pictures of buildings when in an fMRI scanner, the brain response in areas of the brain known to be associated with navigation, such as the parahippocampal gyrus

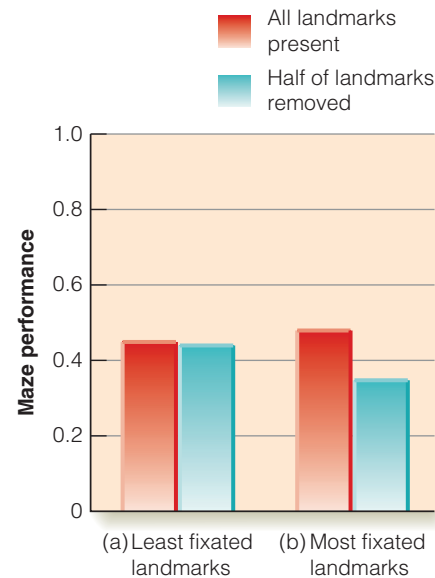


Figure 7.13 Effect of removing landmarks on maze performance. Red = all landmarks are present; green = half have been removed. (a) Removing half of the least fixated landmarks has no effect on performance. (b) Removing half of the most fixated landmarks causes a decrease in performance. Based on Hamid, S. N., Stankiewicz, B., & Hayhoe, M. (2010). Gaze patterns in navigation: Encoding information in large-scale environments. *Journal of Vision*, 10(12):18, 1–11. Figure 4.




Figure 7.14 The human brain, showing three structures important to navigation: the parahippocampal gyrus, the hippocampus, and the retrosplenial cortex. © Cengage Learning 2014

(see **Figure 7.14**), was larger than the response to non-decision-point buildings. Thus, decision-point landmarks are not only more likely to be recognized than non-decision-point landmarks, but they generate greater levels of brain activity.

In another brain scanning experiment, Janzen and van Turennout (2004) had observers first study a film sequence that moved through a “virtual museum” (**Figure 7.15**). Observers were told that they needed to learn their way around the museum well enough to be able to guide a tour

through it. Objects (“exhibits”) were located along the hallway of this museum. Decision-point objects, like the object at (a), marked a place where it was necessary to make a turn. Non-decision-point objects, like the one at (b), were located at a place where a decision was not required.

After studying the museum’s layout in the film, observers were given a recognition test while in an fMRI scanner. They saw objects that had been in the hallway and some objects they had never seen. Their brain activation was measured in the scanner as they indicated whether they remembered seeing each object. **Figure 7.15c** indicates activity in the right parahippocampal gyrus for objects the observers had seen as they learned their way through the museum. The left pair of bars indicates, as we might expect, that for objects that the observers remembered, activation was greater for decision-point objects than for non-decision-point objects. But the most interesting result, indicated by the right pair of bars, was that the advantage for decision-point objects also occurred for objects that were not remembered during the recognition test.

Janzen and van Turennout concluded that the brain automatically distinguishes objects that are used as landmarks to guide navigation. The brain therefore responds not just to the object but also to how relevant that object is for guiding navigation. This means that the next time you are trying to find your way along a route that you have traveled before but aren’t totally confident about, activity in your





Figure 7.15 (a & b) Two locations in the “virtual museum” viewed by Janzen and van Turennout’s (2004) observers. (c) Brain activation during the recognition test for objects that had been located at decision points (red bars) and non-decision points (blue bars). Notice that brain activation was greater for decision-point objects even if they weren’t remembered. Adapted by permission from Macmillan Publishers Ltd., from Janzen, G., & van Turennout, M., Selective neural representation of objects relevant for navigation, *Nature Neuroscience*, 7, 673–677. Copyright 2004.

parahippocampal gyrus may automatically be “highlighting” landmarks that indicate when you should continue going straight, or make a right turn or a left turn, even in cases when you may not remember having seen these landmarks before.

From both the behavioral and physiological experiments we have described, it is apparent that landmarks play an important role in wayfinding. But there is more to wayfinding than landmarks. Before you begin a trip, you need to know which direction to go, and you probably also have a mental “map” of your route and the surrounding area in your mind. You may not think of route planning as involving a map, especially for routes that are very familiar, but research studying people who have lost the ability to find their way because of damage to the brain shows that identifying landmarks is just one of the abilities needed to find one’s way.

The Effect of Brain Damage on Wayfinding A large amount of research shows how the ability to navigate through the environment is affected by damage to various brain structures. We will describe cases that involved damage to two structures that have been shown to be involved in navigation, the retrosplenial cortex and the hippocampus (see Figure 7.14).

Retrosplenial Cortex Damage On the evening of December 11, 2000, a 55-year-old taxi driver was suddenly unable to find his way home from work. He was able to recognize buildings, so he knew where he was, but he couldn’t figure out which direction to turn to get home. He called his wife and got home by following her directions (Ino et al., 2007). When this patient was tested at the hospital, it was found that he had damage to his retrosplenial cortex. Behavioral testing revealed that he could identify buildings and other common objects and was able to remember the positions of objects in a room, but he couldn’t describe or draw routes between his house and familiar places or draw the layout of his house. Results such as these led to the conclusion that this patient had lost his *directional ability*—he couldn’t determine the direction of any familiar destination with respect to his current position, and wasn’t able to use directional information provided by familiar landmarks.

This problem in determining direction is illustrated by another case of retrosplenial cortex damage, a 70-year-old retired schoolteacher who was unable to determine the viewpoints from which photographs of familiar places were taken. For example, the three red arrows in **Figure 7.16** show her judgments of the viewpoint from which she thought a photograph of her garden was taken. These responses were, however, completely different from the correct viewpoint, shown by the green arrow (Suzuki, 1998).

Hippocampus Damage Patient T.T. had been a London taxi driver for 37 years when he contracted a severe case of encephalitis that damaged his hippocampus (Maguire et al., 2006). After the damage, he was unable to find his way around his own neighborhood. T.T. was tested on his ability to drive from one place to another in London by navigating a car in an interactive computer game called “The Getaway,” which accurately depicted




Figure 7.16 Responses of a patient with retrosplenial cortex damage when she was asked to identify the viewpoint of a photograph of her garden. The green arrow indicates the correct viewpoint of the photograph. The three red arrows are the patient’s indications of the viewpoints. She was able to identify the garden table, but she could not indicate the direction from which it was seen.

From Suzuki, K., Yamadori, A., Hayakawa, Y., & Fujii, T. (1998). Pure topographical disorientation related to dysfunction of the viewpoint dependent visual system. *Cortex*, 34, 589–599. Reproduced by permission.

the streets of central London as seen through the front window of a car, including all the buildings and landmarks along the road and some pedestrians as well (**Figure 7.17**).

T.T. was able to do this as well as control subjects, a group of retired London taxi drivers, but only if the route involved just main roads. As soon as it was necessary to navigate along side streets, T.T. became lost, even though he had been taking people on taxi rides through the same side streets for 37 years. Eleanor Maguire and coworkers (2006) concluded that the hippocampus is important for accessing details of routes that were learned long ago.

The research we have described on how the brain is involved in wayfinding has focused on three structures: the parahippocampal gyrus, the retrosplenial cortex, and the hippocampus. Physiological research studying the behavior of patients with



Figure 7.17 A view similar to the one in the video game *The Getaway* (© Sony Computer Entertainment Europe), which duplicates the roadways and buildings of downtown London.

brain damage and analysis of the results of brain scanning experiments have also identified a number of other brain areas involved in various components of wayfinding (Schinazi & Epstein, 2010). The important message of all of these studies, taken together, is that wayfinding is distributed throughout many structures in the brain. This isn't surprising when we consider that wayfinding involves seeing and recognizing objects along a route (perception), paying attention to specific objects (attention), using information stored from past trips through the environment (memory), and combining all this information to create maps that help us relate what we are perceiving to where we are now and where we need to go next.

TEST YOURSELF 7.1

1. What two factors does the ecological approach to perception emphasize?
2. What is optic flow? What are two characteristics of optic flow?
3. What is invariant information? How is invariance related to optic flow?
4. What is observer-produced information? Describe its role in somersaulting and why there is a difference between novices and experts when they close their eyes.
5. Describe the swinging room experiments. What principles do they illustrate?
6. What is the evidence (a) that optic flow provides information for the direction someone is heading and (b) that there are neurons that respond to optic flow?
7. What does research on driving a car and walking tell us about how optic flow may (or may not) be used in navigation? What are some other sources of information for navigation?
8. What is wayfinding? Describe the research of Hamid et al. (computer maze) and Schinazi and Epstein (walking on the Penn campus) that investigated the role of landmarks in wayfinding.
9. What do the brain scanning experiments of Schinazi and Epstein (measuring responses to buildings on the Penn campus) and Janzen and van Turenout (measuring activation when navigating a virtual museum) indicate about brain activity and landmarks?
10. Describe the case studies of patients with damage to their RSP and hippocampus. What conclusions about the function of these structures were reached from these observations?
11. What does it mean to say that wayfinding is "multifaceted"?

Acting on Objects

So far, we have been describing how we move around in the environment. But our actions go beyond walking or driving. One of the major actions we take is reaching to pick something

up, as Serena did on her bike ride when she reached down, grabbed her water bottle, and raised it to her mouth. One of the characteristics of reaching and grasping is that it is usually directed toward a specific object, to accomplish a specific goal. We reach for and grasp a doorknob to open a door; we reach for a hammer to pound a nail. An important concept related to reaching and grasping is *affordances*, which we describe next.

Affordances: What Objects Are Used For

Remember that Gibson's ecological approach involves identifying information in the environment that is useful for perception. Earlier in the chapter we described optic flow, which is created by movement of the observer. Another type of information that Gibson specified is *affordances*—information that indicates what an object is used for. In Gibson's (1979) words, "The affordances of the environment are what it *offers* the animal, what it *provides for* or *furnishes*." A chair, or anything that is sit-on-able, affords sitting; an object of the right size and shape to be grabbed by a person's hand affords grasping; and so on.

What this means is that perception of an object not only includes physical properties, such as shape, size, color, and orientation, that might enable us to recognize the object; our perception also includes information about how the object is used. For example, when you look at a cup, you might receive information indicating that it is "a round white coffee cup, about 5 inches high, with a handle," but your perceptual system would also respond with information indicating "you can pick the cup up" and "you can pour liquid into it." Information such as this goes beyond simply seeing or recognizing the cup; it provides information that can guide our actions toward it. Another way of saying this is that "potential for action" is part of our perception of an object.

One way that affordances have been studied is by looking at the behavior of people with brain damage. Glyn Humphreys and Jane Riddoch (2001) studied affordances by testing patient M.P., who had damage to his temporal lobe that impaired his ability to name objects. M.P. was given a cue, either (1) the name of an object ("cup") or (2) an indication of the object's function ("an item you could drink from"). He was then shown 10 different objects and was told to press a key as soon as he found the object. The results of this testing showed that M.P. identified the object more accurately and rapidly when given the cue that referred to the object's function. Humphreys and Riddoch concluded from this result that M.P. was using his knowledge of an object's affordances to help find it.

Although M.P. wasn't reaching for these objects, it is likely that he would be able to use the information about an object's function to help him take action with respect to the object. In line with this idea, there are other patients with temporal lobe damage who cannot name objects, or even describe how they can be used, but who can pick them up and use them nonetheless.

The Physiology of Reaching and Grasping

An important breakthrough in the study of the physiology of reaching and grasping came with the discovery of ventral (or *what*) and dorsal (or *where/how*) pathways that we described in Chapter 4 (see Figure 4.14).

The Dorsal and Ventral Pathways Remember that D.F., who had damage to her ventral pathway, had difficulty recognizing objects or judging their orientation, but she could “mail” an object by placing it through an oriented opening. The idea that there is one processing stream for perceiving objects and another for acting on them helps us understand what is happening when Serena, sitting at the coffee shop after her ride, reaches for her cup of coffee (**Figure 7.18**). She first identifies the coffee cup among the flowers and other objects on the table (ventral pathway). Once the coffee cup is perceived, she reaches for it, taking into account its location on the table (dorsal pathway). As she reaches, avoiding the flowers, she positions her hand and fingers to grasp the cup (dorsal), taking into account her perception of the cup’s handle (ventral). She then lifts the cup with just the right amount of force (dorsal), taking into account her estimate of how heavy it is based on her perception of its fullness (ventral).

Thus, reaching and picking up a cup involves continually perceiving the position of the cup, shaping the hand and fingers relative to the cup, and calibrating actions in order to accurately grasp the cup and pick it up without spilling any coffee (Goodale, 2011). Even a seemingly simple action like picking up a coffee cup involves a number of areas of the brain, which coordinate their activity to create perceptions and behaviors.

The Parietal Reach Region One of the most important areas of the brain for reaching and grasping is the parietal lobe at the end of the dorsal pathway (Figure 7.7). The areas in the monkey and human parietal cortex that are involved

in reaching for objects have been called the **parietal reach region (PRR)**. This region contains neurons that control not only grasping but also reaching (Connolly et al., 2003). Recently, evidence has been presented suggesting that there are a number of different parietal reach regions in the human parietal lobe (Filimon et al., 2009), and recording from single neurons in a monkey’s parietal lobe has revealed neurons in an area next to the parietal reach region that respond to specific types of hand grips (Fattori et al., 2010).

The procedure for the monkey hand grip experiment, which was carried out by Patrizia Fattori and coworkers (2010), is shown in **Figure 7.19**: (1) The monkey observes a small fixation light in the dark; (2) lights are turned on for half a second to reveal the object to be grasped; (3) the lights go out and then, after a brief pause, the fixation light changes color, signaling that the monkey should reach for the object.

The key part of this sequence occurs when the monkey reaches for the object in the dark. The monkey knows what the object is from seeing it when the lights were on (a round ball in this example), so while it is reaching for the object in the dark, it adjusts its grip to match the object. A number of different objects were used, as shown in **Figure 7.19b**, each of which required a different grip. **AVL**

The key result of the experiment is that there are neurons that respond best to specific grips. For example, neuron A in **Figure 7.20** responds best to “whole hand prehension” whereas neuron B responds best to “advanced precision grip.” There are also neurons, like C, that respond to a number of different grips. Remember that when these neurons were firing, the monkey was reaching for the object in the dark, so the firing reflected not perception but the monkey’s actions.

In a follow-up experiment on the same monkeys, Fattori and coworkers (2012) discovered neurons that responded not only when a monkey was preparing to grasp a specific object, but also when the monkey *viewed* that specific object. An example of this type of neuron, which Fattori calls **visuomotor grip cells**, is a neuron that initially responds when the monkey sees a specific object, and then also responds as the monkey is



Figure 7.18 Picking up a cup of coffee: (a) perceiving and recognizing the cup, (b) reaching for it, and (c) grasping and picking it up. This action involves coordination between perceiving and action that is carried out by two separate streams in the brain, as described in the text. From Goldstein, E. B., *Cognitive Psychology*, 3rd ed. © 2011 Wadsworth, a part of Cengage Learning, Inc. Reproduced by permission. www.cengage.com/permissions.




Figure 7.19 (a) The monkey's task in Fattori and coworkers' (2010) experiment. The monkey always looks at the small light above the sphere. The monkey sees the object to be grasped when the lights go on, then reaches for and grasps the object once the lights go off and the fixation light changes color. (b) Four of the objects used in the task. Each one involves a different type of grasping movement. Based on Fattori, P., Raos, V., Breveglieri, R., Bosco, A., Marzocchi, N., & Galletti, C. (2010). The dorsomedial pathway is not just for reaching: Grasping neurons in the medial parieto-occipital cortex of the Macaque monkey. *Journal of Neuroscience*, 30, 342–349. Figure 2b, c.




Figure 7.20 Results of Fattori and coworkers' (2010) experiment showing how three different neurons respond to reaching and grasping each of the objects. Neuron A responds best to “whole hand prehension” (starred record). Neuron B responds to “advanced precision grip.” Neuron C responds to all of the grips. Based on Fattori, P., Raos, V., Breveglieri, R., Bosco, A., Marzocchi, N., & Galletti, C. (2010). The dorsomedial pathway is not just for reaching: Grasping neurons in the medial parieto-occipital cortex of the Macaque monkey. *Journal of Neuroscience*, 30, 342–349. Figure 2.

forming its hand to grasp the same object. This type of neuron is therefore involved in both perception (identifying the object by seeing) and action (reaching for the object and gripping it with the hand).

Avoiding Other Objects When Reaching When we reach, we have to take into account not only the location toward which we are reaching, so we can direct our hand toward that location, but also the location of other nearby objects, so we can avoid them as we reach. Serena faced this problem when she had to reach toward her coffee cup while avoiding the vase of flowers and the glass of orange juice.

The fact that obstacle avoidance is also controlled by the parietal regions responsible for reaching was demonstrated in an experiment by Igor Schindler and coworkers (2004), who tested two patients with parietal lobe damage who had trouble pointing to visual stimuli, a condition called **optic ataxia**. These ataxia patients and a group of normal control subjects were presented with two cylinders, separated by 8 to 10 inches (**Figure 7.21a**). Their task was to reach between the two cylinders and touch anywhere on a gray strip located 20 cm behind the cylinders. The cylinders were moved to different positions, as shown by the top views of pairs of cylinders in **Figure 7.21b**.

The arrows indicate where the subject's hand passed between the cylinders as he or she reached to touch the strip. Notice that the control subjects (red arrows) changed their reach in response to changes in the cylinders' position, shifting their reach to the left when the cylinders were shifted to the left. In contrast, the reach of the ataxia patients was the same for all arrangements of the cylinders, as shown for one of the patients by the blue arrows. In other words, they didn't take account of the varying locations of the obstacles. Schindler concludes from this result that the dorsal stream, which was damaged in the ataxia patients, not only provides guidance as we reach toward an object but also guides us away from potential obstacles.




Figure 7.21 (a) Subjects in Schindler and coworkers' (2004) experiment had to reach between the two cylinders to touch a gray strip located behind the cylinders. (b) The pairs of cylinders in Schindler and coworkers' (2004) experiment were located in different positions on different trials, as shown in this top view. The red arrows show that control subjects adjusted their reach to compensate for the different locations of the cylinders. The blue arrows, which show the data for one of the ataxia patients, indicate that the patients' reach stayed the same for all arrangements of the cylinders. Based on Schindler, I., Rice, N. J., McIntosh, R. D., Rossetti, Y., Vighetto, A., & Milner, D.A. (2004). Automatic avoidance of obstacles is a dorsal stream function: Evidence from optic ataxia. *Nature Neuroscience*, 7, 779–784.

Observing Other People's Actions

We not only take action ourselves, but we regularly watch other people take action. This “watching others act” is most obvious when we watch other people's actions on TV or in a movie, but it also occurs any time we are around someone else who is doing something. One of the most exciting outcomes of research studying the link between perception and action was the discovery of neurons in the premotor cortex (**Figure 7.7**) called *mirror neurons*.

Mirroring Others' Actions in the Brain

In the early 1990s, Giacomo Rizzolatti and coworkers (2006; also see di Pellegrino et al., 1992; Gallese et al., 1996) were investigating how neurons in the monkey's premotor cortex fired as the monkey performed actions like picking up a toy or a piece of food. Their goal was to determine how neurons fired as the monkey carried out specific actions. But as sometimes happens in science, they observed something they didn't expect. When one of the experimenters picked up a piece of food while the monkey was watching, neurons in the monkey's cortex fired. What was so unexpected was that the neurons that fired to observing the experimenter pick up the food were the same ones that had fired earlier when the monkey had itself picked up the food.

This initial observation, followed by many additional experiments, led to the discovery of **mirror neurons**—neurons that respond both when a monkey observes someone else grasping an object such as food on a tray (**Figure 7.22a**) and when the monkey itself grasps the food (**Figure 7.22b**; Rizzolatti et al., 2006). They are called mirror neurons because




Figure 7.22 Response of a mirror neuron. (a) Response to watching the experimenter grasp food on the tray. (b) Response when the monkey grasps the food. (c) Response to watching the experimenter pick up food with a pair of pliers. Reprinted from Rizzolatti, G., et al., Premotor cortex and the recognition of motor actions, *Cognitive Brain Research*, 3, 131–141. Copyright 2000, with permission from Elsevier.

the neuron’s response to watching the experimenter grasp an object is similar to the response that would occur if the monkey were performing the same action. Just looking at the food causes no response, and watching the experimenter grasp the food with a pair of pliers, as in **Figure 7.22c**, causes only a small response (Gallese et al., 1996; Rizzolatti et al., 2000).

Most mirror neurons are specialized to respond to only one type of action, such as grasping or placing an object somewhere. Although you might think that the monkey may have been responding to the anticipation of receiving food, the type of object made little difference. The neurons responded just as well when the monkey observed the experimenter pick up an object that was not food. **VL**

But could the mirror neurons simply be responding to the pattern of motion? The fact that the neuron does not respond when watching the experimenter pick up the food with pliers argues against this idea. Further evidence that mirror neurons are doing more than just responding to a particular pattern of motion is the discovery of neurons that respond to sounds that are *associated with* actions. These neurons in the premotor cortex, called **audiovisual mirror neurons**, respond when a monkey performs a hand action *and* when it hears the sound associated with this action (Kohler et al., 2002). For example, the results in **Figure 7.23** show the response of a neuron that fires (a) when the monkey sees and hears the experimenter break a peanut, (b) when the monkey just sees the experimenter break the peanut, (c) when the monkey just hears the sound of the breaking peanut, and (d) when the *monkey* breaks the peanut. What this means is that just *hearing* a peanut breaking or just *seeing* a peanut being broken causes activity that is also associated with the perceiver’s *action* of breaking a peanut. These neurons are responding, therefore, to what is “happening”—breaking a peanut—rather than to a specific pattern of movement.

Predicting People’s Intentions

Some researchers have proposed that there are mirror neurons that respond not just to *what* is happening but to *why* something is happening, or more specifically, to the *intention* behind what is happening. To understand what this means, let’s return to Serena in the coffee shop. As we see her reach

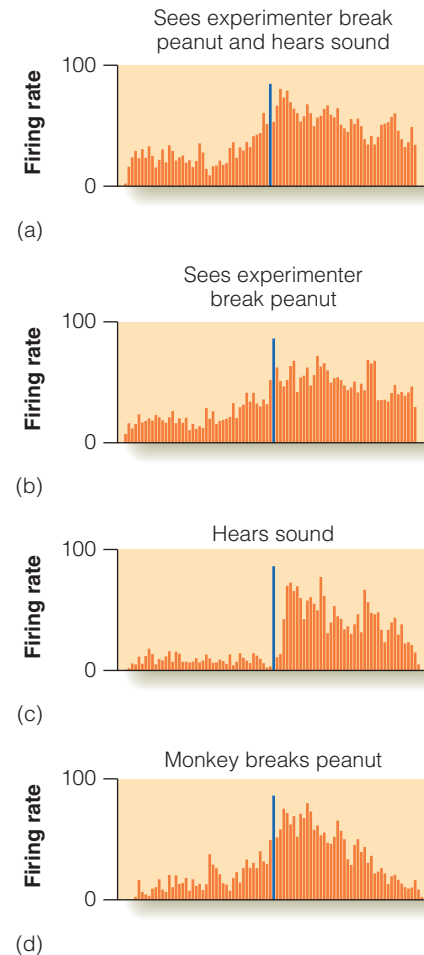


Figure 7.23 Response of an audiovisual mirror neuron to four different stimuli. From Kohler, E., et al., 2002, Hearing sounds, understanding actions: Action representation in mirror neurons. *Science*, 297, 846–848. Copyright © 2002 by AAAS. Reprinted with permission from AAAS.

for her coffee cup, we might wonder why she is reaching for it. One obvious answer is that she intends to drink some coffee, although if we notice that the cup is empty, we might instead decide that she is going to take the cup back to the counter to get a refill, or if we know that she never drinks more than one cup, we might decide that she is going to place the cup in the used cup bin. Thus, there are a number of different intentions that may be associated with the same action.




Figure 7.24 Images from the Context, Action, and Intention film clips viewed by Iacoboni and coworkers' (2005) subjects. See text for details.

From Iacoboni, M., Molnar-Szakacs, I., Gallese, V., Buccino, G., Mazziotta, J. C., & Rizzolatti, G. (2005). Grasping the intentions of others with one's own mirror neuron system. *PLoS Biology*, 3(3), e79. Used by permission.

What is the evidence that the response of mirror neurons can be influenced by different intentions? Mario Iacoboni and coworkers (2005) provide this evidence in an experiment in which they measured subjects' brain activity as they watched short film clips represented by the stills in **Figure 7.24**. Stills for the two Intention films, on the right, show a hand reaching in to pick up a cup, but there is an important difference between the two scenes. In the top panel, the table is neatly set up, the food is untouched, and the cup is full of tea. In the bottom panel, the table is a mess, the food has been eaten, and the cup appears to be empty. Iacoboni hypothesizes that it is likely that viewing the top film would lead the viewer to infer that the person picking up the cup intends to drink from it, and that viewing the bottom film would lead the viewer to infer that the person is cleaning up.

Iacoboni's subjects also viewed the control films shown in the other panels. The Context film showed the table setting, and the Action film showed the hand reaching in to pick up an isolated cup. The reason these two types of films were presented was that they contained the visual elements of the intention films, but didn't suggest a particular intention.

When Iacoboni compared the brain activity in the Intention films to the activity in the Context and Action films, he found that the Intention films caused greater activity than the control films in areas of the brain known to have mirror neuron properties. **Figure 7.25** shows that the amount of activity was least in the Action condition, was higher for the Cleaning Up condition, and was highest for the Drinking condition. Based on the increased activity for the two Intention conditions, Iacoboni concluded that the mirror neuron area is involved with understanding the intentions behind the actions shown in the films. He reasoned that if the mirror neurons were just signaling the action of picking up the cup, then a similar response would occur regardless of whether a context surrounding the cup was present. Mirror neurons,

according to Iacoboni, code the "why" of actions and respond differently to different intentions.

If mirror neurons do, in fact, signal intentions, how do they do it? One possibility is that the response of these neurons is determined by the chain of motor activities that could be *expected* to happen in a particular context (Fogassi et al., 2005; Gallese, 2007). For example, when a person picks up a cup with the intention of drinking, the next expected actions would be to bring the cup to the mouth and then to drink some coffee. However, if the intention is to clean up, the expected action might be to carry the cup over to the sink. According to this idea, mirror neurons that respond to different intentions are responding to the action that is happening *plus* the sequence of actions that is most likely to follow, given the context.

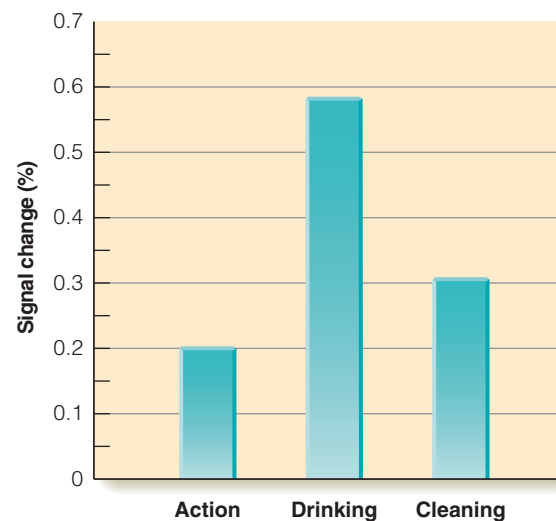


Figure 7.25 Iacoboni and coworkers' (2005) results, showing the brain response for the Action, Drinking, and Cleaning conditions.

© Cengage Learning 2014

The exact functions of mirror neurons in humans are still being actively researched (Caggiano et al., 2011; de Lange et al., 2008; Gazzola et al., 2007; Kilner, 2011). In addition to proposing that mirror neurons signal what is happening as well as the intentions behind various actions, researchers have also proposed that mirror neurons help us understand (1) communications based on facial expressions (Buccino et al., 2004; Ferrari et al., 2003); (2) gestures used while speaking (Gallese, 2007); (3) the meanings of sentences (Gallese, 2007); and (4) differences between ourselves and others (Uddin et al., 2007). As might be expected from this list, it has also been proposed that mirror neurons play an important role in guiding social interactions (Rizzolatti & Sinigaglia, 2010; Yoshida et al., 2011).

As with any newly discovered phenomenon, more research is needed before we can state with more certainty exactly what the function of mirror neurons is. Consider that when feature detectors that respond to oriented moving lines were discovered in the 1960s, some researchers proposed that these feature detectors could explain how we perceive objects. With the information available at the time, this was a reasonable proposal. However, later, when neurons that respond to faces, places, and bodies were discovered, researchers revised their initial proposals to take these new findings into account. In all likelihood, a similar process will occur for mirror neurons. Some of the proposed functions will be confirmed, but others may need to be revised. This evolution of thinking about what research results mean is a basic property not only of research in perception but of scientific research in general.

SOMETHING TO CONSIDER:

Action-Based Accounts of Perception

The traditional approach to perception has focused on how the environment is *represented* in the nervous system and in the perceiver's mind. According to this idea, the purpose of visual perception is to create a representation in the mind of whatever we are looking at. Thus, if you look at a scene and see buildings, trees, grass, and some people, your perception of the buildings, trees, grass, and people is representing what is "out there," and so accomplishes vision's purpose of representing the environment.

But as you might have suspected after reading this chapter, many researchers believe that the purpose of vision is not to create a representation of what is out there but to guide our actions. We can appreciate the reasoning behind this idea by imagining a situation in which action is important for survival. Consider a monkey foraging for food in the forest. The monkey's color perception enables it to see some orange fruit that stands out against green leaves. The monkey reaches for the fruit and eats it. Of course, seeing (and perhaps smelling)

the fruit is crucial, because it makes the monkey aware that the fruit is present. But the second step—reaching for the fruit—is just as important, because the monkey can't live on visual experiences alone. It has to reach for and grab the fruit in order to survive.

Although there may be situations—such as looking at paintings in an art gallery or looking out at a misty lake in the morning—when seeing what is out there is an end in itself, the vast majority of our experience involves a two-step process: first *perceiving* an object or scene and then *taking action* toward the objects or within the scene.

The idea that action is crucial for survival has been described by Mel Goodale (2011) as follows: "Many researchers now understand that brains evolved not to enable us to think (or perceive), but to enable us to move and interact with the world" (p. 17). According to this idea, perception may provide valuable information about the environment, but taking a step beyond perception and acting on this information enables us to survive so we can perceive another day (Milner & Goodale, 2006).

The idea that the purpose of perception is to enable us to interact with the environment has been taken a step further by researchers who have turned the equation around from "action depends on perception" to "perception depends on action" or "people perceive their environment in terms of their ability to act on it." This last statement, by Jessica Witt (2011), is based on the results of many experiments, some of which involve sports. For example, Witt and Dennis Proffitt (2005) presented a series of circles to softball players just after they had finished a game and asked them to pick the circle that best corresponded to the size of a softball. When they compared the players' estimates to their batting averages from the just-completed game, they found that batters who hit well perceived the ball to be bigger than batters who were less successful. **VL**

Other experiments that have focused on sports have shown that tennis players who have recently won report that the net is lower (Witt & Sugovic, 2010), and that subjects who were most successful at kicking football field goals estimated the goal posts to be farther apart (Witt & Dorsch, 2009). The field goal experiment is especially interesting because the effect occurred only after they had attempted 10 field goals. Before they began, the estimates of the poor kickers and the good kickers were the same.

The sports examples all involved making judgments after doing either well or poorly. This supports the idea that perception can be affected by performance. What about situations in which the person hasn't carried out any action but has an expectation about how difficult it would be to perform that action? For example, what if people who were physically fit and people who were not physically fit were asked to estimate the steepness of a hill? When Mukul Bhalla and Dennis Proffitt (1999) asked people ranging from varsity athletes to people who didn't work out regularly to estimate the slant of steep hills, they found that the least fit people (as measured by heart rate and oxygen consumption during and

after exercise) judged the hills as being steeper. The reason for this, according to Bhalla and Proffitt, is that over time people's general fitness level affects their perception of how difficult it will be to carry out various types of physical activity, and this in turn affects their perception of these activities. Thus, a person who isn't very fit experiences steep hills as being difficult to climb, and this causes them to perceive the hills as being steeper even if they are just looking at them (Proffitt, 2009).

The idea that the expected difficulty of carrying out an action can influence a person's judgment of an object's properties was also studied by Adam Doerrfeld and coworkers (2011), who asked subjects to estimate the weight of a basket of golf balls before and after lifting the basket. Subjects made this estimate under two conditions: (1) solo, in which the subject expected that he or she would be lifting the basket alone, and (2) joint, in which the subject expected that he or she would be lifting the basket with another person. The actual weight of the basket of golf balls was 20 pounds. Before lifting the basket, the subjects estimated that the basket weighed 21 pounds if they thought they would be lifting it alone, and 17.5 pounds if they thought they would be lifting it with another person. After lifting the basket, the average estimate was about 20 pounds for both conditions. Doerrfeld and coworkers conclude from this result that anticipation of how difficult a task will be can influence the perception of an object's properties.

There are, however, researchers who question whether the perceptual judgments measured in some of the experiments we have described are actually measuring perception. Subjects might be affected, they suggest, by "judgmental bias," caused by their expectations about what they think will happen in a particular situation. For example, Bhalla and Proffitt (1999), who found that people who were not in good physical condition judged hills as being steeper, also found that people who were wearing a heavy backpack judged hills to be steeper. Bhalla and Proffitt interpreted this result as showing that wearing the heavy backpack influenced the person's perception of steepness. An alternative interpretation is that perhaps the subjects' *expectation* that hills could appear steeper when carrying something heavy might cause them to *say* a hill appears steeper when they are wearing a heavy backpack, even though their *perception* of the hill's steepness was actually not affected (Durgin et al., 2010; Loomis & Philbeck, 2008; Woods et al., 2009).

This explanation highlights a basic problem in measuring perception in general: Our measurement of perception is based on people's responses, and there is no guarantee that these responses accurately reflect what a person is perceiving. Thus, as pointed out above, there may be some instances in which subjects' responses may reflect not what they are perceiving, but what they think they should be perceiving. Even though some experiments may be open to criticism (Durgin et al., 2010; Proffitt, 2009), it is important to note that there are some experiments that do demonstrate a

relationship between a person's ability to act and perception (Creem-Regehr & Kunz, 2010).

The results of the experiments demonstrating this relationship between ability to act and perception are consistent with J. J. Gibson's idea of affordances, described earlier (page 163). Affordances, according to Gibson, are an object's "possibilities for action." Thus, perception of a particular object is determined both by what the object looks like and by the way we might interact with it.

This brings us to the following statement by J. J. Gibson, from his final book, *The Ecological Approach to Perception* (1979): "Perceiving is an achievement of the individual, not an appearance in the theater of his consciousness. It is a keeping-in-touch with the world, an experiencing of things, rather than a having of experiences" (p. 239). This statement did not lead to much research when it was proposed, but years later many researchers have embraced the idea that perception is not just "an appearance in the theater of consciousness," but is the first step toward taking action in the environment. In addition, some researchers have gone a step farther and suggested that action, or the potential for action, may affect perception.

TEST YOURSELF 7.2

1. What is an affordance? Describe the results of the experiments on patient M.P. that illustrates the operation of affordances.
2. Describe the early experiments that showed that there are neurons in the parietal cortex that respond to goal-directed reaching.
3. How does the idea of *what* (ventral) and *how* (dorsal) streams help us describe an action such as reaching for a coffee cup?
4. Describe Fattori et al.'s experiments on "grasping neurons" and "visuomotor grip cells."
5. What is the parietal reach region?
6. Describe the experiment on optic ataxia patients that shows that the dorsal stream is involved in helping to avoid obstacles.
7. What are mirror neurons? What is the evidence that mirror neurons aren't just responding to a specific pattern of motion?
8. Describe Iacoboni's experiment that suggested that there are mirror neurons that respond to intentions.
9. What is a possible mechanism that might be involved in mirror neurons that respond to intentions?
10. What are some of the proposed functions of mirror neurons? What is the scientific status of these functions?
11. Describe the action-based account of perception. In your discussion, indicate (a) why some researchers think the brain evolved to enable us to take action; (b) how experiments have demonstrated a link between perception and "ability to act."

THINK ABOUT IT

1. It is a common observation that people tend to slow down as they are driving through long tunnels. Explain the possible role of optic flow in this situation. (p. 154)
2. We have seen that gymnasts appear to take visual information into account as they are in the act of executing a somersault. In the sport of synchronized diving, two people execute a dive simultaneously from two side-by-side diving boards. They are judged based on how well they execute the dive and how well the two divers are synchronized with each other. What environmental stimuli do you think synchronized divers need to take into account in order to be successful? (p. 155)
3. Can you identify specific environmental information that you use to help you carry out actions in the environment? This question is often particularly relevant to athletes.
4. If mirror neurons do signal intentions, what does that say about the role of top-down and bottom-up processing in determining the response of mirror neurons? (p. 166)
5. How do you think the response of your mirror neurons might be affected by how well you know a person whose actions you were observing? (p. 166)
6. How does your experience in interacting with the environment (climbing hills, playing sports) correspond or not correspond to the findings of the “potential for action” experiments described in the Something to Consider section? (p. 169)

KEY TERMS

Affordance (p. 163)

Audiovisual mirror neuron (p. 167)

Ecological approach to perception (p. 154)

Focus of expansion (FOE) (p. 154)

Gradient of flow (p. 154)

Invariant information (p. 155)

Landmarks (p. 160)

Mirror neuron (p. 166)

Optic ataxia (p. 166)

Optic flow (p. 154)

Parietal reach region (PRR) (p. 164)

Self-produced information (p. 155)

Visual direction strategy (p. 159)

Visuomotor grip cells (p. 164)

Wayfinding (p. 160)

MEDIA RESOURCES

CourseMate

Go to CengageBrain.com to access Psychology CourseMate, where you will find the Virtual Labs plus an interactive eBook, flashcards, quizzes, videos, and more.

Virtual Labs

The Virtual Labs are designed to help you get the most out of this course. The Virtual Lab icons direct you to specific media demonstrations and experiments designed to help you visualize what you are reading about. The numbers below indicate the number of the Virtual Lab you can access through Psychology CourseMate.

7.1 Flow From Moving Down a Hallway (p. 154)

A computer-generated program showing the optic flow that occurs when moving through a patterned hallway. (Courtesy of William Warren)

7.2 Optic Flow Over Surface (p. 154)

Flow from moving across a texture field. (Courtesy of Zhi Li and Frank Durgin)

7.3 Stimuli Used in Warren’s Experiment (p. 157)

Moving stimulus pattern seen by observers in William Warren’s experiment. (Courtesy of William Warren)

7.4 Eye Movements While Driving (p. 159)

Eye tracking while driving under different conditions. (Courtesy of Farid Kandil)

7.5 Optic Flow and the Visual Control of Locomotion (p. 159)

A review of optic flow and visual direction as sources of information for locomotion, and description of ongoing research. (Courtesy of William Warren)

7.6 Blind Walking Experiment (p. 159)

Subjects carrying out instructions in a blind walking experiment. (Courtesy of John Philbeck)

7.7 Monkey Grasping (p. 164)

Shows how monkeys grasped objects in the Fattori and colleagues’ (2010) experiment. (Courtesy of Patrizia Fattori)

7.8 Monkey Perception While Grasping (p. 164)

Shows what monkey sees during the Fattori grasping experiment. (Courtesy of Patrizia Fattori)

7.9 Mirror Neurons (p. 167)

Describes EEG research of Jaime Pineda that is related to mirror neurons in humans.

7.10 Connection Between Action and Perception in Tennis (p. 169)

Description of the results of Jessica Witt's experiments on the connection between tennis performance and how tennis players perceive the speed of the ball and the height of the net. (Courtesy of Karin Heineman, American Institute of Physics)



Ashley Cooper/Corbis

Perceiving Motion

CHAPTER CONTENTS

Functions of Motion Perception

Motion Helps Us Understand Events in Our Environment
Motion Attracts Attention
Motion Provides Information About Objects

Studying Motion Perception

When Do We Perceive Motion?
Comparing Real and Apparent Motion
What We Want to Explain

Motion Perception: Information in the Environment

Motion Perception: Retina/Eye Information

The Reichardt Detector
Corollary Discharge Theory

Motion Perception and the Brain

The Movement Area of the Brain
Effect of Lesioning and Microstimulation
Motion From a Single Neuron's Point of View

Motion and the Human Body

Apparent Motion of the Body
Motion of Point-Light Walkers

Representational Momentum: Motion Responses to Still Pictures

SOMETHING TO CONSIDER: Event Perception

Think About It

◀ Our perception of motion depends on the movement of images across our retina, as would occur if these birds flew across our field of view; on signals generated by movement of our eyes, which would occur if we followed the birds' movement; and on cognitive mechanisms based on what we have learned by observing our environment. Our perception of this picture as birds in motion is based on our general knowledge of birds and on cues to motion such as the blurred images of some of the birds.

VL The Virtual Lab icons direct you to specific animations and videos designed to help you visualize what you are reading about. Virtual Labs are listed at the end of the chapter, keyed to the page on which they appear, and can be accessed through Psychology CourseMate.

Some Questions We Will Consider:

- Why do some animals freeze in place when they sense danger? (p. 177)
- When we scan or walk through a room, the image of the room moves across the retina, but we perceive the room and the objects in it as remaining stationary. Why does this occur? (p. 182)
- Why is motion of the human body “special”? (p. 190)

We are always taking action, either dramatically—as in Serena's bike ride in Chapter 7 (page 153)—or routinely, as in reaching for a coffee cup or walking across a room. Whatever form action takes, it involves motion, and one of the things that makes the study of motion perception both fascinating and challenging is that we are not simply passive observers of the motion of others. We are often moving ourselves. Thus, we perceive motion when we are stationary, as when we are watching other people cross the street (**Figure 8.1a**), and we also perceive motion as we ourselves are moving, as might happen when playing basketball (**Figure 8.1b**). We will see in this chapter that both the “simple” case of a stationary observer perceiving motion and the more complicated case of a moving observer perceiving motion involve complex “behind-the-scenes” mechanisms.



Cathrine Wessel/Corbis

(a)



George Doyle/Stockbyte/Getty Images

(b)

Figure 8.1 Motion perception occurs (a) when a stationary observer perceives moving stimuli, such as this couple crossing the street; and (b) when a moving observer, like this basketball player, perceives moving stimuli, such as the other players on the court.

Functions of Motion Perception

Motion perception has a number of different functions, ranging from providing us with updates about what is happening to helping us perceive things such as the shapes of objects and people's moods. Perhaps most important of all, especially for animals, the perception of motion is intimately linked to survival.

Motion Helps Us Understand Events in Our Environment

As you walk through a shopping mall, looking at the displays in the store windows, you are also observing other actions—a group of people engaged in an animated conversation, a salesperson rearranging piles of clothing and then walking over to the cash register to help a customer, a TV program in a restaurant that you recognize as a dramatic moment in a soap opera.

Much of what you observe involves information provided by motion. The gestures of the people in the group indicate the intensity of their conversation; the motions of the salesperson indicate what she is doing and changes in motion indicate when she has shifted to a new task; and motion indicates, even in the absence of sound, that something

important is happening in the soap opera (Zacks, 2004; Zacks & Swallow, 2007).

A particularly compelling demonstration of motion's power to indicate what is happening was provided by Fritz Heider and Marianne Simmel (1944), who showed a 2½-minute animated film to subjects and asked them to describe what was happening in the movie. The movie consisted of a “house” and three “characters”—a small circle, a small triangle, and a large triangle. These three geometric objects moved around both inside and outside the house, and sometimes interacted with each other (**Figure 8.2**).




Figure 8.2 One image from the film used by Heider and Simmel (1944). The objects moved in various ways, going in and out of the “house” and sometimes interacting with each other. The nature of the movements led subjects to make up stories that often described the objects as having feelings, motivations, and personalities. Adapted from Heider, F., & Simmel, M. (1944). An experimental study of apparent behavior. *American Journal of Psychology*, 13, 243–259.

Although the characters were geometric objects, the subjects created stories to explain the objects' actions, and often gave them humanlike characteristics and personalities. For example, one account described the small triangle and circle as a couple who were trying to be alone in the house when the big triangle ("a bully") entered the house and interrupted them. The small triangle didn't appreciate this intrusion and attacked the big triangle. Who would have thought the world of geometric objects could be so exciting?

Returning to the world of people, motion perception is essential for our ability to move through the environment. As we saw in Chapter 7 when we described how people navigate (see page 157), one source of information about where we are going and how fast we are moving is the way objects in the environment flow past us as we move. As a person moves forward, objects move relative to the person in the opposite direction. This movement, called *optic flow* (Figures 7.1 and 7.2, page 154), provides information about the walker's direction and speed. In Chapter 7 we discussed how we can use this information to help us stay on course.

While motion provides information about what is going on and where we are moving, it provides information for more subtle actions as well. Consider, for example, the action of pouring water into a glass. As we pour the water, we watch the level rise, and this helps us know when to stop pouring. We can appreciate the importance of this ability by considering the case of a 43-year-old woman who lost the ability to perceive motion when she suffered a stroke that damaged an area of her cortex involved in motion perception. Her condition, called *akinetopsia* (blindness to motion), made it difficult for her to pour tea or coffee into a cup because the liquid appeared frozen, so she couldn't perceive the fluid rising in the cup and had trouble knowing when to stop pouring. It was also difficult for her to follow dialogue because she couldn't see the motions of a speaker's face and mouth (Zihl et al., 1983, 1991).

But the most disturbing effect of her brain damage occurred when people suddenly appeared or disappeared, because she couldn't see them approaching. Crossing the street presented serious problems because at first a car might seem far away, but then suddenly, without warning, it would appear very near. This disability was not just a social inconvenience but enough of a threat to the woman's well-being that she rarely ventured outside into the world of moving—and sometimes dangerous—objects. This case of a breakdown in the ability to perceive motion provides a dramatic demonstration of the importance of motion perception in day-to-day life.

Motion Attracts Attention

As you try to find your friend among the sea of faces in the student section of the stadium, you realize that you have no idea where to look. But you suddenly see a person waving

and recognize that it is your friend. The ability of motion to attract attention is called **attentional capture**. This effect occurs not only when you are consciously looking for something but also while you are paying attention to something else. For example, as you are having a conversation, your attention may suddenly be captured by something moving in your peripheral vision.

The fact that movement can attract attention plays an important role in animal survival. You have probably seen animals freeze in place when they sense danger. If a mouse's goal is to avoid being detected by a cat, one thing it can do is to stop moving. Freezing in place not only eliminates the attention-attracting effects of movement, it also makes it harder for the cat to differentiate between the mouse and its surroundings.

Motion Provides Information About Objects

The idea that not moving can help an animal blend into the background is illustrated by the following demonstration. **VL**

DEMONSTRATION

Perceiving a Camouflaged Bird

For this demonstration, you will need to prepare stimuli by photocopying the bird and the hatched-line pattern in **Figure 8.3**. Then cut out the bird and the hatched pattern so



Figure 8.3 The bird becomes camouflaged when the random lines are superimposed on it. When the bird is moved relative to the lines, it becomes visible, an example of how movement enhances the perception of form. From Regan, D. (1986). Luminance contrast: Vernier discrimination. *Spatial Vision*, 1, 305–318. Reprinted by permission of David Regan.

they are separated. Hold the picture of the bird up against a window during the day. Turn the copy of the hatched pattern over so the pattern is facing out the window (the white side of the paper should be facing you) and place it over the bird. If the window is adequately illuminated by daylight, you should be able to see the hatched pattern. Notice how the presence of the hatched pattern makes it more difficult to see the bird. Then, slide the bird back and forth under the pattern, and notice what happens to your perception of the bird (from Regan, 1986).

The stationary bird is difficult to see when it is covered by the pattern because the bird and the pattern are made up of similar lines. But as soon as all the elements of the bird begin moving in the same direction, the bird becomes visible. Movement has perceptually organized all the elements of the bird, so they create a figure that is separated from the background. Returning to our mouse hiding from the cat, we can say that it is to the mouse's advantage to freeze because this decreases the chances that the mouse will become perceptually separated from its surroundings in the cat's mind. **VL**

You might say, in reaction to the camouflaged bird demonstration, that although motion does make the bird easy to perceive amid the tangle of obscuring lines, this seems like a special case because most of the objects we see are not camouflaged. But if you remember our discussion from Chapter 5 (page 97) about how even clearly visible objects may be ambiguous, you can appreciate how motion of an object can reveal characteristics of the object that might not be obvious from a single, stationary view (**Figure 8.4a**). Movement of an observer around an object causes a similar effect: viewing the "horse" in **Figure 8.4b**

from different perspectives reveals that its shape is not exactly what you may have expected based on your initial view. Thus, our own motion relative to objects is constantly adding to the information we have about the objects, and most relevant to this chapter, we receive similar information when objects move relative to us. Observers perceive shapes more rapidly and accurately when an object is moving (Wexler et al., 2001).

Studying Motion Perception

To describe how motion perception is studied, the first question we will consider is: When do we perceive motion?


When Do We Perceive Motion?

The answer to this question may seem obvious: We perceive motion when something moves across our field of view. Actual motion of an object is called **real motion**. Perceiving a car driving by, people walking, or a bug scurrying across a tabletop are all examples of the perception of real motion.

There are also a number of ways to produce the perception of motion that involve stimuli that are not moving. Perception of motion when there actually is none is called **illusory motion**. The most famous, and best studied, type of illusory motion is called **apparent motion**. We introduced apparent motion in Chapter 5 when we told the story of Max Wertheimer's observation that when two stimuli in



Figure 8.4 (a) The shape and features of this car are revealed as different aspects of it become visible as it moves. (b) Moving around this "horse" reveals its true shape.



© Cengage Learning 2014



Bruce Goldstein

Figure 8.5 Apparent motion (a) between two dots when they are flashed one after the other; (b) on a moving sign. Our perception of words moving across a lighted display is so compelling that it is often difficult to realize that signs like this one are simply dots flashing on and off.

slightly different locations are alternated with the correct timing, an observer perceives one stimulus moving back and forth smoothly between the two locations (**Figure 8.5a**). This perception is called **apparent motion** because there is no actual (or real) motion between the stimuli. This is the basis for the motion we perceive in movies, on television, and in moving signs that are used for advertising and entertainment (**Figure 8.5b**). (Also see Figure 5.15, page 101.)

Induced motion occurs when motion of one object (usually a large one) causes a nearby stationary object (usually smaller) to appear to move. For example, the moon usually appears stationary in the sky. However, if clouds are moving past the moon on a windy night, the moon may appear to be racing through the clouds. In this case, movement of the larger object (clouds covering a large area) makes the smaller, but actually stationary, moon appear to be moving.

Motion aftereffects occur when viewing a moving stimulus for 30 to 60 seconds causes a stationary stimulus to appear to move. One example of a motion aftereffect is the **waterfall illusion** (Addams, 1834) (**Figure 8.6a**). If you look at a waterfall for 30 to 60 seconds (be sure it fills up only part of your field of view) and then look off to the side at part of the scene that is stationary, you will see everything you are looking at—rocks, trees, grass—appear to move up for a few seconds (**Figure 8.6b**).

Researchers studying motion perception have investigated all the types of perceived motion described above—and a number of others as well (Blaser & Sperling, 2008; Cavanagh, 2011). Our purpose, however, is not



(a)



(b)

Bruce Goldstein

Figure 8.6 The waterfall movement aftereffect. (a) Observation of motion in one direction, such as occurs when viewing a waterfall, can cause (b) the perception of motion in the opposite direction, indicated by the arrows, when viewing stationary objects in the environment.

to understand every type of motion perception but to understand some of the principles governing motion perception in general. To do this, we will focus on real and apparent motion.

Comparing Real and Apparent Motion

For many years, researchers treated the apparent motion created by flashing stationary objects or pictures and the real motion created by actual motion through space as though they were separate phenomena, governed by different mechanisms. However, there is ample evidence that these two types of motion have much in common. For example, Axel Larsen and coworkers (2006) presented three types of displays to a person in an fMRI scanner: (1) a *control condition*, in which two dots in slightly different positions were flashed simultaneously (**Figure 8.7a**); (2) a *real motion display*, in which a small dot moved back and forth (**Figure 8.7b**); and (3) an *apparent motion display*, in which dots were flashed one after another so that they appeared to move back and forth (**Figure 8.7c**). **VL**

Larsen's results are shown below the dot displays. The blue-colored area in Figure 8.7a is the area of visual cortex activated by the control dots, which are perceived as two dots simultaneously flashing on and off with no motion between them. Each dot activates a separate area of the cortex. In Figure 8.7b, the red indicates the area of cortex activated by real movement of the dot. In Figure 8.7c, the yellow indicates the area of cortex activated by the apparent motion display. Notice that the activation associated with apparent motion is similar to the activation for the real motion display. Two flashed dots that result in apparent motion activate the area of brain representing the space between the positions of the flashing dots even though no stimulus was presented there.

Because of the similarities between the neural responses to real and apparent motion, researchers study both types of motion together and concentrate on discovering general

mechanisms that apply to both. In this chapter, we will follow this approach as we look for general mechanisms of motion perception.

What We Want to Explain

Our goal is to understand how we perceive things that are moving. At first this may seem like an easy problem. For example, **Figure 8.8a** shows what Maria sees when she looks straight ahead as Jeremy walks by. Because she doesn't move her eyes, Jeremy's image sweeps across her retina. Explaining motion perception in this case seems straightforward because as Jeremy's image moves across Maria's retina, it stimulates a series of receptors one after another, and this stimulation signals Jeremy's motion.

Figure 8.8b shows what Maria sees when she follows Jeremy's motion with her eyes. In this case, Jeremy's image remains stationary on Maria's foveas as he walks by. This adds an interesting complication to explaining motion perception, because although Jeremy's image remains stationary on her retina, Maria perceives Jeremy as moving. This means that motion perception can't be explained just by the motion of an image across the retina.

Let's consider what happens if Jeremy isn't present, and Maria scans the room by moving her eyes from left to right. When Maria does this, the images of the walls and objects in the room move to the left across her retina (**Figure 8.8c**), but Maria doesn't see the room or its contents as moving. In this case, there is motion across the retina but no perception that objects are moving. This is another example of why we can't simply consider what is happening on the retina. **Table 8.1** summarizes the three situations in Figure 8.8.

In the sections that follow, we will consider a number of different approaches to explaining motion perception, with the goal being to explain each of the situations in Figure 8.8 and Table 8.1. We begin by considering an approach that focuses on how information in the environment signals motion.





Figure 8.7 Three conditions in Larsen's (2006) experiment: (a) control condition; (b) real motion; (c) apparent motion (flashing dots). Stimuli are shown on top, and the resulting brain activation is shown below. In (c), the brain is activated in the space that represents the area between the two dots, where movement was perceived but no stimuli were present. From Larsen, A., Madsen, K. H., Lund, T. E., & Bundesen, C., Images of illusory motion in primary visual cortex. *Journal of Cognitive Neuroscience*, 18, 1174–1180. © 2006 by the Massachusetts Institute of Technology.



(a) Jeremy walks past Maria; Maria's eyes are stationary (creates local disturbance in optic array)



(b) Jeremy walks past Maria; Maria follows him with her eyes (creates local disturbance in optic array)



(c) Scans scene by moving her eyes from left to right (creates global optic flow)

Figure 8.8 Three motion situations: (a) Maria is stationary and looks straight ahead as Jeremy walks past; (b) Maria follows Jeremy's movement with her eyes; (c) Maria scans the room by moving her eyes to the right. (The optic array and optic flow are described in the next section.)

© Cengage Learning

TABLE 8.1 Conditions for Perceiving and Not Perceiving Motion Depicted in Figure 8.8

	SITUATION	OBJECT	EYES	IMAGE ON OBSERVER'S RETINA	OBJECT MOVEMENT PERCEIVED?
1	Look straight as an object moves past	Moves	Stationary	Moves	YES
2	Follow a moving object with eyes	Moves	Move	Stationary	YES
3	Look around the room	Stationary	Move	Moves	NO

© Cengage Learning

Motion Perception: Information in the Environment

From the three situations in Figure 8.8, we saw that motion perception can't be explained by considering just what is happening on the retina. A solution to this problem was suggested by J. J. Gibson, who founded the ecological approach to perception. In Chapter 7 we noted that Gibson's approach (1950, 1966, 1979) involves looking for information in the environment that is useful for perception (see page 154). This information for perception, according to Gibson, is located not on the retina but "out there" in the environment. He thought about information in the environment in terms of the **optic array**—the structure created by the surfaces, textures, and contours of the environment—and he focused on how movement of the observer causes changes in the optic array. Let's see how this works by returning to Jeremy and Maria in Figure 8.8.

In Figure 8.8a, when Jeremy walks across Maria's field of view, portions of the optic array become covered as he walks by and then are uncovered as he moves on. This result is called a **local disturbance in the optic array**. This local disturbance in the optic array occurs when Jeremy moves relative to the environment, covering and uncovering the stationary background. According to Gibson, this local disturbance in the optic array provides information that Jeremy is moving relative to the environment.

In Figure 8.8b, Maria follows Jeremy with her eyes. Remember that Gibson doesn't care what is happening on the retina. Even though Jeremy's image is stationary on the retina, the same local disturbance information that was available when Maria was keeping her eyes still—Jeremy covering and uncovering parts of the array—remains available when she is moving her eyes, and this local disturbance information indicates that Jeremy is moving.

However, when Maria scans the scene in Figure 8.8c, something different happens: As her eyes move across the scene from left to right, everything around her—the walls, the window, the trash can, the clock, and the furniture—moves to the left of her field of view. A similar situation would occur if Maria were to walk through the scene. The fact that everything moves at once in response to movement of the observer's eyes or body is called **global optic flow**; this signals that the environment is stationary. Thus, according to Gibson, motion is perceived when one part of the visual scene moves relative to the rest of the scene, and no motion is perceived when the entire field moves, or remains stationary.

Motion Perception: Retina/ Eye Information

Gibson's approach focuses on information that is "out there" in the environment. Another approach to explaining the various movement situations in Figure 8.8 is to consider the neural signals that travel from the eye to the brain.

The Reichardt Detector

An early neural explanation for motion perception is a neural circuit proposed by Werner Reichardt (1969) called the **Reichardt detector**, which results in neurons that fire to movement in one direction. **Figure 8.9** illustrates the basic principle of the Reichardt detector. Excitation and inhibition are arranged so that movement in one direction creates inhibition that eliminates neural responding, whereas movement in the opposite direction creates excitation that enhances neural responding.

We can understand how this works by following what happens as a spot of light moves across the retinal receptors. Figures 8.9a and b show what happens when the light is moving from left to right. Receptor A is stimulated first. The synapse between receptor A and E is excitatory (indicated by the Y), so stimulation of A excites E (indicated by green). Receptor E makes an inhibitory synapse with F (indicated by the vertical line), so F is inhibited (indicated by orange). While this is occurring, the light has moved to the right to receptor B and causes it to respond and to send an excitatory signal to F, but since F has already been inhibited by E, it does not fire (**Figure 8.9b**). Thus, when the light is moving to the right, the




Figure 8.9 Reichardt circuit. Green indicates excitation; orange indicates inhibition. (a) and (b) When the receptors are stimulated from left to right, neuron I does not fire. (c) and (d) When the receptors are stimulated from right to left, neuron I fires. © Cengage Learning 2014

signals from receptors A and B do not get past F and therefore never reach I, the neuron at the end of the circuit. This process is repeated as the stimulus moves across the remaining receptors. The net result is that when the light is moving across the receptors from left to right, neuron I does not respond.

The outcome is different, however, when the light starts at receptor D and moves to the left. Receptor D sends a signal to H, which causes it to fire (**Figure 8.9c**) and to excite neuron I. When the light moves to the left and stimulates receptor C, it activates neuron G, which sends inhibition back to H. This inhibition, however, arrives too late, because H has already fired and has stimulated neuron I (**Figure 8.9d**). This process is repeated as the stimulus moves across the remaining receptors. Thus, when the light is moving to the left, the inhibition arrives too late to stop the signal from getting to neuron I, so neuron I fires. Neuron I, therefore, does not fire to movement to the right (Figure 8.9a and b) but does fire to movement to the left (Figure 8.9c and d).

Corollary Discharge Theory

Reichardt detectors can detect motion in a specific direction, but they can only explain the situation in Figure 8.8a, when an image (in this case, the image of Jeremy) sweeps across the receptors. In order to explain situations like those in Figure 8.8b (when Maria moves her eyes to follow Jeremy's movements) and Figure 8.8c (when Maria scans the room), we need to take into account not only how the image is moving on the retina but also how the eye is moving. **Corollary discharge theory** takes eye movements into account. The first step in understanding corollary discharge theory is to consider how neural signals associated with the retina and with the eye muscles are related to the three situations in Figure 8.8.

Signals From the Retina and the Eye Muscles Corollary discharge theory explains motion perception by taking into account the following signals, which are generated by movement of a stimulus on the retina and by movement of the eyes.

1. An **image displacement signal (IDS)** (**Figure 8.10a**) occurs when an image moves across receptors in the retina, as when Jeremy walks across Maria's field of view while she stares straight ahead.
2. A **motor signal (MS)** (**Figure 8.10b**) occurs when a signal is sent from the brain to the eye muscles. This signal occurs when Maria moves her eyes to follow Jeremy as he walks across the room.
3. A **corollary discharge signal (CDS)** is a copy of the motor signal that, instead of going to the eye muscles, is sent to a different place in the brain (Figure 8.10b). This is analogous to using the "cc" (copy) function when sending an email message. The email goes to the person it is addressed to, and a copy of the email is simultaneously sent to someone else at another address.

Now that we have introduced these signals, we can see a solution to our problem by asking what situations 1 and 2,

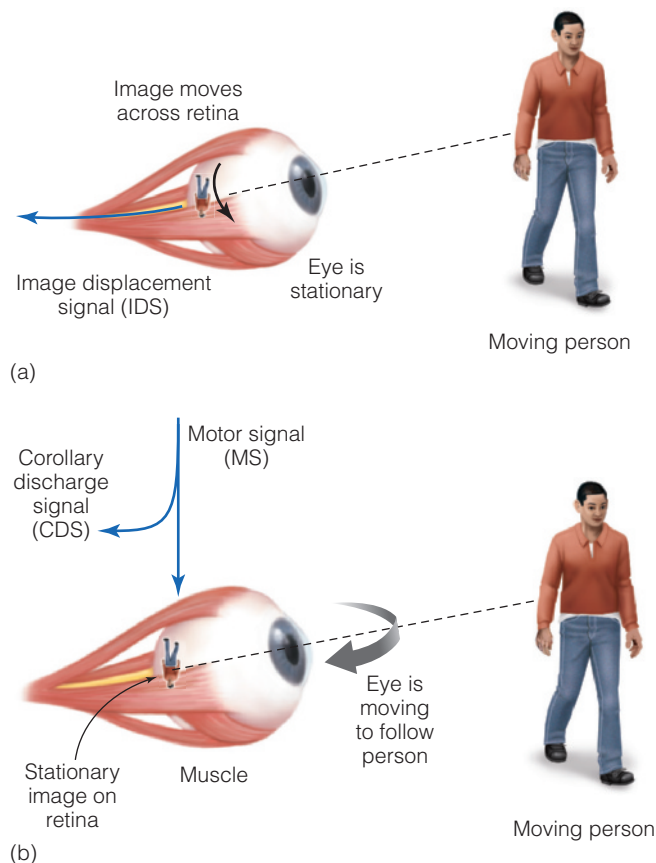


Figure 8.10 (a) When the image of an object moves across the retina, movement of the image across the retina creates an image displacement signal (IDS). (b) When a motor signal (MS) to move the eyes is sent to the eye muscles, so the eye can follow a moving object, there is a corollary discharge signal (CDS), which splits off from the motor signal. © Cengage Learning

in which the object is perceived to move, have in common. We can answer that question by focusing on the two signals that are transmitted toward the brain: the image displacement signal (IDS) and the corollary discharge signal (CDS). In situation 1, when Maria keeps her eyes stationary and Jeremy's image moves across her retina, only an IDS occurs. In situation 2, in which Maria moves her eyes to follow Jeremy so Jeremy's image doesn't move across her retina, only a CDS occurs. So perhaps the solution is this: When only one type of signal, either the IDS or the CDS, is sent to the brain, motion is perceived. Furthermore, if both signals occur, as happens in situation 3, when an observer scans the room as in Figure 8.8c, then no motion is perceived. This solution is, in fact, the basis of corollary discharge theory.

According to corollary discharge theory, the brain contains a structure or mechanism called the **comparator** that receives both the IDS and the CDS. The operation of the comparator is governed by the rules illustrated in **Figure 8.11**. If just one type of signal reaches the comparator—either the IDS (**Figure 8.11a**) or the CDS (**Figure 8.11b**)—it relays a message to the brain that “movement has occurred,” and motion is perceived. But if both the CDS and IDS reach the comparator at the same time (**Figure 8.11c**), they cancel each






Figure 8.11 According to corollary discharge theory, (a) when the IDS reaches the comparator alone, a signal is sent to the brain and motion is perceived; (b) when the CDS reaches the comparator alone, a signal is sent to the brain and motion is perceived; (c) if both a CDS and an IDS reach the comparator simultaneously, they cancel each other, so no signals are sent to the brain and no motion is perceived.

© Cengage Learning

other, so no signal is sent to the area of the brain responsible for motion perception. This handles our problem, because motion is perceived in situations 1 and 2, in which only one type of signal is present, but isn't perceived in situation 3, when both types of signal are present. **VL**

Upon hearing this explanation, students often ask where the comparator is located. The answer is that the comparator is most likely not located in one specific place in the brain but may involve a number of different structures. Similarly, the CDS probably originates from a number of different places in the brain (Sommer & Crapse, 2010; Sommer & Wurtz, 2008). The important thing for our purposes is that corollary discharge theory proposes that the visual system takes into account both information about stimulation of the receptors and information about movement of the eyes. And although we can't pinpoint exactly where the CDS and comparator are located, there is evidence that supports the theory. Here is some of the behavioral and physiological evidence.

Behavioral Evidence for Corollary Discharge Theory

These two demonstrations create a perception of motion even though there is no motion across the retina.



Figure 8.12 Afterimage stimulus. © Cengage Learning

DEMONSTRATION

Eliminating the Image Displacement Signal With an Afterimage

Illuminate the circle in **Figure 8.12** with your desk lamp and look at it for about 60 seconds. Then go into your closet (or a completely dark room) and observe what happens to the circle's afterimage (blink to make it come back if it fades) as you look around. Notice that the afterimage moves in synchrony with your eye motions (**Figure 8.13**).




Figure 8.13 Afterimage demonstration. When the eye moves in the dark, the image remains stationary (the bleached area on the retina indicated by the red oval), but a CDS is sent to the comparator, so the afterimage appears to move. © Cengage Learning

Why does the afterimage appear to move when you move your eyes? The answer cannot be that an image is moving across your retina because the circle's image always remains at the same place on the retina. (The circle's image on the retina has created a circular area of bleached visual pigment, which remains in the same place no matter where the eye is looking.) Without motion of the stimulus across the retina, there is no image displacement signal. However, the motor signals sent to move your eyes are creating a corollary discharge signal, which reaches the comparator alone, so the afterimage appears to move (**Figure 8.11b**).

DEMONSTRATION

Seeing Motion by Pushing on Your Eyelid

Pick a point in the environment and keep looking at it while *very gently* pushing back and forth on the side of your eyelid, as shown in **Figure 8.14**. As you do this, you will see the scene move.



Bruce Goldstein

Figure 8.14 Why is this woman smiling? Because when she pushes on her eyelid, while keeping her eye fixed on one place, she sees the world jiggle.

Why do you see motion when you push on your eyelid? Lawrence Stark and Bruce Bridgeman (1983) did an experiment in which they instructed observers to keep looking at a particular point while pushing on their eyelid. Because the observers were paying strict attention to the instructions (“Keep looking at that point!”), the push in their eyelid didn’t cause their eyes to move. This lack of movement occurred because the observer’s eye muscles were pushing back against the force of the finger to keep the eye in place. According to corollary discharge theory, the motor signal sent to the eye muscles to hold the eye in place created a corollary discharge signal, which reached the comparator alone, as in Figure 8.11b, so Stark and Bridgeman’s observers saw the scene move (also see Bridgeman

& Stark, 1991; Ilg, Bridgeman, & Hoffmann, 1989). (See “Think About It” #3 on page 196 for a question related to this explanation.)

These demonstrations support the central idea proposed by corollary discharge theory that there is a signal (the corollary discharge) that indicates when the observer moves, or tries to move, his or her eyes. When the theory was first proposed, there was little physiological evidence to support it, but now there is a great deal of physiological evidence for the theory.

Physiological Evidence for Corollary Discharge Theory

In both of our demonstrations, there was a corollary discharge signal but no image displacement signal. What would happen if there were no corollary discharge signal but there *was* an image displacement signal? That is apparently what happened to R.W., a 35-year-old male who experienced vertigo (dizziness) anytime he moved his eyes or experienced motion when he looked out the window of a moving car.

A brain scan revealed that R.W. had lesions in an area of his cortex called the medial superior temporal (MST) area (refer back to Figure 7.7). Behavioral testing of R.W. also revealed that as he moved his eyes, the stationary environment appeared to move with a velocity that matched the velocity with which he was moving his eyes (Haarmeier et al., 1997). Thus, when he moved his eyes to the left, there was an IDS, because images were moving across his retina to the right but the damage to his brain had apparently eliminated the CDS. Because only the IDS reached the comparator, R.W. saw motion when there actually was none.

Other physiological evidence for the theory comes from experiments that involve recording from neurons in the monkey’s cortex. **Figure 8.15** shows the response recorded from a motion-sensitive neuron in the monkey’s extrastriate cortex.

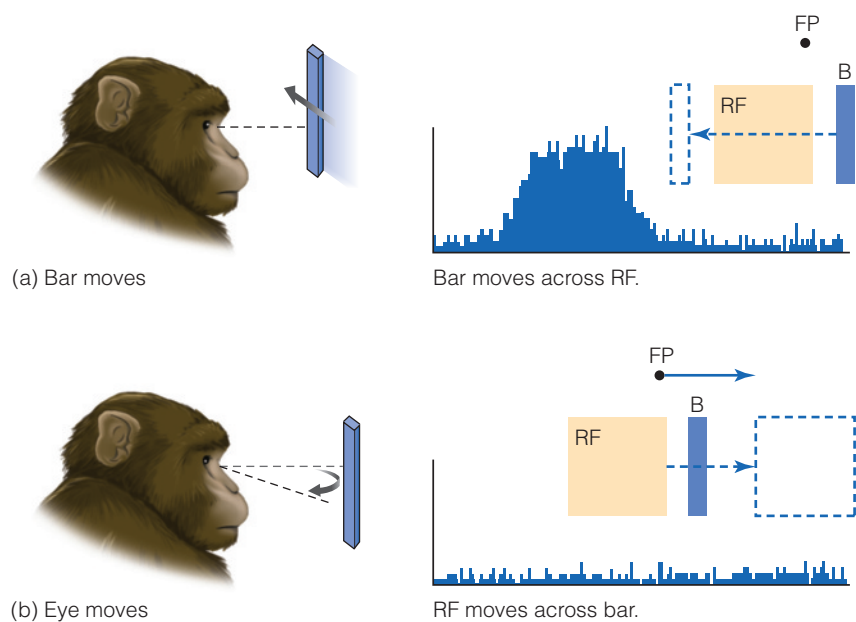


Figure 8.15 Responses of a real-motion neuron in the extrastriate cortex of a monkey. In both cases, a bar (B) sweeps across the neuron’s receptive field (RF) as the monkey looks at a fixation point (FP). (a) The neuron fires when the bar moves to the left across the receptive field. (b) The neuron doesn’t fire when the eye moves to the right even though this also causes the bar to move across the receptive field. Adapted from Galletti, C., & Fattori, P. (2003). Neuronal mechanisms for detection of motion in the field of view. *Neuropsychologia*, 41, 1717–1727.

This neuron responds strongly when the monkey looks steadily at the fixation point (FP) as a moving bar sweeps across the neuron's receptive field (RF) (Figure 8.15a). But what if the monkey moves its eyes to follow a moving fixation point so its eyes sweep across a stationary bar (Figure 8.15b)? In this case, the bar's image will sweep across the neuron's receptive field, just as it did in Figure 8.15a. Even though the bar is sweeping across the receptive field, just as before, the neuron doesn't fire (Galletti & Fattori, 2003).

This neuron is called a **real-motion neuron** because it responds only when the stimulus moves and doesn't respond when the eye moves, even though the stimulus on the retina—a bar sweeping across the cell's receptive field—is the same in both situations. This real-motion neuron must be receiving information like the corollary discharge signal, which tells the neuron when the eye is moving. Real-motion neurons have also been observed in many other areas of the cortex (Battaglini et al., 1996; Robinson & Wurtz, 1976), and more recent research has begun to determine where the corollary discharge signal is acting in the brain (Sommer & Wurtz, 2006; Wang et al., 2007).

TEST YOURSELF 8.1

1. Describe four different functions of motion perception.
2. Describe four different situations that can result in motion perception. Which of these situations involve real motion, and which involve illusions of motion?
3. What is the evidence for similar neural responding to real motion and apparent motion?
4. Describe Gibson's ecological approach to motion perception. What is the advantage of this approach? (Give a specific example of how the ecological approach can explain the situations in Figure 8.8b and c.)
5. Describe the operation of the neural circuit that creates the Reichardt detector.
6. Describe the corollary discharge model. In your description, indicate (1) what the model is designed to explain; (2) the three types of signals—image displacement signal, motor signal, corollary discharge signal; and (3) when these signals do and do not cause motion perception when reaching the comparator.




Figure 8.16 Moving dot displays used by Newsome, Britten, and Movshon (1989). These pictures represent moving dot displays that were created by a computer. Each dot survives for a brief interval (20–30 microseconds), after which it disappears and is replaced by another randomly placed dot. Coherence is the percentage of dots moving in the same direction at any point in time. (a) Coherence = 0 percent; (b) Coherence = 50 percent; (c) Coherence = 100 percent. From Newsome, W. T., & Paré, E. B. (1988).

A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *Journal of Neuroscience*, 8, 2201–2211. Reproduced by permission.

Motion Perception and the Brain

In this section we will focus on the brain, and specifically on the middle temporal (MT) area, and the medial superior temporal (MST) area, both of which play important roles in the perception of motion.

The Movement Area of the Brain

When we described Hubel and Wiesel's (1959, 1965) pioneering work on receptive fields, we saw that they recorded from neurons in the visual receiving area that responded to bars that moved in a specific direction (see Figure 3.29). Another area that contains many directionally sensitive cells is the middle temporal (MT) area. Evidence that the MT cortex is specialized for processing information about motion is provided by experiments that have used moving dot displays in which the direction of motion of individual dots can be varied.

Figure 8.16a represents a display in which all of the dots are moving in random directions. William Newsome and coworkers (1995) used the term **coherence** to indicate the degree to which the dots move in the same direction. When the dots are all moving in random directions, coherence is 0 percent. Figure 8.16b represents a coherence of 50 percent, as indicated by the darkened dots, which means that at any point in time half of the dots are moving in the same direction. Figure 8.16c represents 100 percent coherence, which means that all of the dots are moving in the same direction.

Newsome and coworkers used these moving dot stimuli to determine the relationship between (1) a monkey's ability to judge the direction in which dots were moving and (2) the response of a neuron in the monkey's MT cortex. They found that as the dots' coherence increased, two things happened: (1) the monkey judged the direction of motion more accurately, and (2) the MT neuron fired more rapidly. The monkey's behavior and the firing of the MT neurons were so closely related that the researchers could predict one from the other. For example, when the dots' coherence was




Figure 8.17 The perceptual cycle from Chapter 1. Newsome measured the physiology–perception relationship by simultaneously recording from neurons and measuring the monkey’s behavioral response. Other research we have discussed, such as Hubel and Wiesel’s receptive field studies, have measured the stimulus–physiology relationship.
© Cengage Learning 2014

0.8 percent, the monkey was not able to judge the direction of the dots’ motion and the neuron’s response did not differ appreciably from its baseline firing rate. But at a coherence of 12.8 percent—so, out of 200 moving dots, about 25 were moving in the same direction—the monkey judged the direction of the dots that were moving together correctly on virtually every trial, and the MT neuron always fired faster than its baseline rate.

We can appreciate the importance of Newsome’s experiments by considering the following three basic relationships in **Figure 8.17**, which we introduced in Chapter 1 (see Figure 1.10):

- *The stimulus–perception relationship (green arrow):* Presenting a stimulus and determining whether motion is perceived. For example, when an object moves fast enough, we perceive movement; when an array of dots are moving in the same direction, we perceive movement in that direction.
- *The stimulus–physiology relationship (orange arrow):* Presenting a movement stimulus and measuring neural responding. For example, in the experiment shown in Figure 8.15a, a moving bar caused a response in a monkey’s cortex.
- *The physiology–perception relationship (red arrow):* Measuring the relationship between physiological responding and perception. This is the relationship measured by Newsome and coworkers because they measured the response of the MT neurons to the moving dots and also measured the monkey’s perception of the moving dots.

The simultaneous measurement of neural firing and perception indicated by the red arrow is extremely difficult because before the recording experiments can begin, monkeys must be trained for months to indicate the direction in which they perceive the dots moving. (They are given a reward when they correctly signal the direction of movement.) Only after this extensive behavioral training can the monkey’s perception and neural firing be measured simultaneously. The payoff, however, is that the relationship between physiology and perception is measured *directly*, thereby completing the triangle by providing the third relationship in Figure 8.17.

Effect of Lesioning and Microstimulation

Measuring perception and the firing of neurons in the monkey’s MT cortex simultaneously is one way of showing that the MT cortex is important for motion perception. The role of the MT cortex has also been studied by determining how the perception of motion is affected by (1) lesioning (destroying or deactivating) some or all of the MT cortex or (2) electrically stimulating neurons in the MT cortex.

A monkey with an intact MT cortex can begin detecting the direction dots are moving when coherence is as low as 1 to 2 percent. However, after the MT is lesioned, the coherence must be 10 to 20 percent before monkeys can begin detecting the direction of motion (Newsome & Paré, 1988; also see Movshon & Newsome, 1992; Newsome et al., 1995; Pasternak & Merigan, 1994). This example of the physiology–perception relationship provides further evidence linking the firing of MT neurons to the perception of the direction of motion.

Another way this link between the MT cortex and motion perception has been studied is by electrically stimulating neurons in the MT cortex using a technique called *microstimulation*.

METHOD Microstimulation

Microstimulation is achieved by lowering a small wire electrode into the cortex and passing a weak electrical charge through the tip of the electrode. This weak shock stimulates neurons that are near the electrode tip and causes them to fire, just as they would if they were being stimulated by chemical neurotransmitters released from other neurons.

Remember from Chapter 4 that neurons are organized in columns in the cortex, with neurons in the same column responding best to one orientation (page 80). Because neurons that respond to a specific direction of movement are also organized into columns, it is possible to activate neurons that respond to a specific direction of motion by applying microstimulation to a particular column.

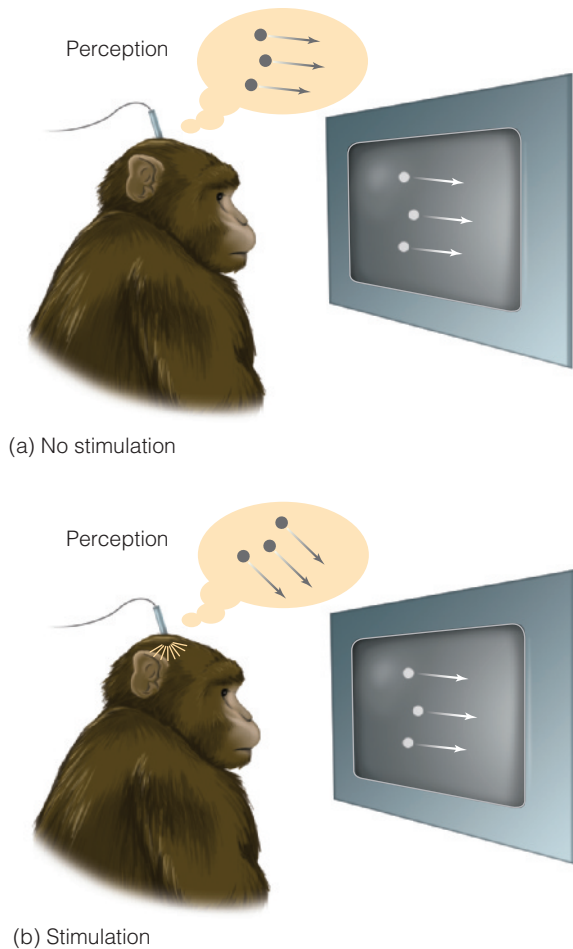


Figure 8.18 (a) A monkey judges the motion of dots moving horizontally to the right. (b) When a column of neurons that prefer downward motion is stimulated, the monkey judges the same motion as being downward and to the right. © Cengage Learning

Anthony Movshon and William Newsome (1992) used this microstimulation procedure in an experiment in which a monkey was looking at dots moving in a particular direction while indicating the direction of motion it was perceiving. For example, **Figure 8.18a** shows that as the monkey observed dots moving to the right, it reported that the dots were moving to the right. But **Figure 8.18b** shows that when Movshon and Newsome stimulated a column of MT neurons that preferred downward motion, the monkey began responding as though the dots were moving downward and to the right. The fact that stimulating the MT neurons shifted the monkey's perception of the direction of movement provides more evidence linking MT neurons and motion perception.

In addition to the MT cortex, another area involved in motion perception is the nearby medial superior temporal (MST) area (see Figure 7.7). But motion activates other areas as well. Remember from Chapter 4 that there are areas specialized to respond to faces (the fusiform face area) and bodies (the extrastriate body area), yet these objects also activate many other areas of the brain (Figure 4.22). Similarly,

the MT and MST cortex are specialized to respond to motion, yet motion also activates a number of other areas distributed across the brain (Fischer et al., 2012).

Motion From a Single Neuron's Point of View

Having established that the MT cortex is specialized for perceiving motion, we will now look at a close-up of how motion perception is served by the firing of single neurons within the MT cortex. The obvious answer to the question of how the firing of neurons can signal the direction in which an object is moving is that as an image of the object sweeps across the retina, it activates directionally selective neurons that respond to movement in a specific direction (see Figure 3.29).

Although this appears to be a straightforward solution to signaling the direction an object is moving, it turns out that the response of individual directionally selective neurons does not provide sufficient information to indicate the direction of movement. We can understand why this is so by considering how a directionally selective neuron would respond to movement of a vertically oriented pole like the one being carried by the woman in **Figure 8.19**.




Figure 8.19 The pole's overall motion is horizontally to the right (blue arrows). The ellipse represents the area in an observer's field of view that corresponds to the receptive field of a cortical neuron on the observer's retina. The pole's motion across the receptive field is also horizontal to the right (red arrows). © Cengage Learning

We are going to focus on the pole, which is essentially a vertical bar. The ellipse represents the area of the receptive field of a neuron in the cortex that responds when a vertical bar moves to the right across the neuron's receptive field. Figure 8.19 shows the pole entering the receptive field on the left. As the pole moves to the right, it moves across the receptive field in the direction indicated by the red arrow, and the neuron fires.

But what happens if the woman climbs some steps? **Figure 8.20** shows that as she walks up the steps, she and the pole are now moving up and to the right (blue arrow). We know this because we can see the woman and the flag moving up. But the neuron, which only sees movement through the narrow view of its receptive field, only receives information about the rightward movement. You can demonstrate this for yourself by doing the following demonstration.




Figure 8.20 In this situation, the pole's overall motion is up and to the right (blue arrows). However, the pole's motion across the receptive field is horizontal to the right (red arrows), as in Figure 8.19. Thus, the receptive field "sees" the same motion for motion that is horizontal and motion that is up and to the right. © Cengage Learning




Figure 8.21 Moving a pencil behind an aperture in the "Movement of a Bar Across an Aperture" demonstration. See text for details.

© Cengage Learning

DEMONSTRATION

Movement of a Bar Across an Aperture

Make a small aperture, about 1 inch in diameter, by creating a circle with the fingers of your left hand, as shown in **Figure 8.21** (or you can create a circle by cutting a hole in a piece of paper). Then orient a pencil vertically, and move the pencil from left to right behind the circle, as in **Figure 8.21a**. As you do this, focus on the direction that the *front edge* of the pencil appears to be moving across the aperture. Now, again holding the pencil vertically, position the pencil below the circle, as shown in **Figure 8.21b**, and move it up behind the aperture at a 45-degree angle (being careful to keep its orientation vertical). Again, notice the direction in which the *front edge* of the pencil appears to be moving across the aperture.

If you were able to focus only on what was happening inside the aperture, you probably noticed that the direction that the front edge of the pencil was moving appeared the same whether the pencil was moving (a) horizontally to the right or (b) up and to the right. In both cases, the front edge of the pencil moves across the aperture horizontally, as indicated by the red arrow. Another way to state this is that the movement of an edge across an aperture occurs *perpendicular to the direction in which the edge is oriented*. Because the pencil in our demonstration was oriented vertically, motion through the aperture was horizontal.

Because the motion of the edge was the same in both situations, a single directionally selective neuron would fire similarly in (a) and (b), so based just on the activity of this neuron, it isn't possible to tell whether the pencil is moving horizontally to the right or upward at an angle. The fact that viewing only a small portion of a larger stimulus can result in misleading information about the direction in which the stimulus is moving is called the **aperture problem**.

The visual system appears to solve the aperture problem by pooling the responses of a number of neurons. Evidence that the MT cortex may be involved in pooling the responses from a number of neurons was provided by an experiment by Christopher Pack and Richard Born (2001), in which they determined how neurons in the monkey's MT cortex responded to moving oriented lines like the pole or our pencil. They found that the MT neurons' initial response to the stimulus, about 70 msec after the stimulus was presented, was determined by the orientation of the bar. Thus the neurons responded in the same way to a vertical bar moving horizontally to the right and a vertical bar moving up and to the right (red arrows in Figure 8.21). However, 140 ms after presentation of the moving bars, the neurons began responding to the *actual* direction in which the bars were moving (blue arrows in Figure 8.21). Apparently, MT neurons receive signals from a number of neurons in the striate cortex and then combine these signals to determine the actual direction of motion.

Can you think of another way a neuron might indicate that the pole in Figure 8.20 is moving up and to the right? One of my students tried the demonstration in Figure 8.21 and noticed that when he followed the directions for the demonstration, the edge of the pencil did appear to be moving horizontally across the aperture, whether the pencil was moving horizontally or up at an angle. However, when he moved the pencil so that he could see its tip moving through the aperture, as in **Figure 8.22**, he could tell that the pencil was moving up. Thus, a neuron could use information about the end of a moving object (such as the tip of the pencil) to determine its direction of motion. As it turns out, neurons that could signal this information, because they respond to the ends of moving objects, have been found in the striate cortex (Pack et al., 2003).




Figure 8.22 The circle represents a neuron's receptive field. When the pencil is moved up and to the right, as shown, movement of the tip of the pencil provides information indicating that the pencil is moving up and to the right. © Cengage Learning

What all of this means is that the “simple” situation of an object moving across the visual field as an observer looks straight ahead is not so simple because of the aperture problem. The visual system apparently solves this problem (1) by using information from neurons in the MT cortex that pool the responses of a number of directionally selective neurons, and (2) by using information from neurons in the striate cortex that respond to the movement of the ends of objects (also see Rust et al., 2006; Smith et al., 2005; Zhang & Britten, 2006). **VL**

Motion and the Human Body

We have just seen that experiments using dots and lines as stimuli have taught us a great deal about the mechanisms of motion perception. But what about the more complex stimuli created by moving humans and animals that are so prevalent in our environment? We will now consider two examples of the ways in which researchers have studied how we perceive movement of the human body.

Apparent Motion of the Body

Earlier in this chapter we described *apparent motion* as the perception of motion that occurs when two stimuli that are in slightly different locations are presented one after the other. Even though these stimuli are stationary, movement is perceived back and forth between them if they are alternated with the correct timing. Generally, this movement follows a principle called the **shortest path constraint**—apparent movement tends to occur along the shortest path between two stimuli.

Maggie Shiffrar and Jennifer Freyd (1990, 1993) had observers view photographs like the ones in **Figure 8.23a**, with the photographs alternating rapidly. Notice that in the first picture, the woman's hand is in front of her head, and in the second, it is behind her head. According to the shortest path constraint, motion should be perceived in a straight line between the hands in the alternating photos, which means observers would see the woman's hand as moving through her head, as shown in **Figure 8.23b**. This is, in fact, exactly what happens when the pictures are alternated very rapidly (five or more times a second), even though motion through the head is physically impossible. **VL**

While the straight-line motion of the hand through the head is an interesting result, the most important result occurred when the rate of alternation was slowed. When the pictures were alternated less than five times per second, observers began perceiving the motion as shown in **Figure 8.23c**, so the hand appeared to move around the woman's head. These results are interesting for two reasons: (1) They show that the visual system needs time to process information in order to




Figure 8.23 The two pictures in (a) are photographs similar to those used in Shiffrar and Freyd’s (1993) experiment. The pictures were alternated either rapidly or more slowly. (b) When alternated rapidly, observers perceived the hand as moving through the head. (c) When alternated more slowly, the hand was seen as moving around the head.

perceive the movement of complex meaningful stimuli; and (2) they suggest that there may be something special about the meaning of the stimulus—in this case, the human body—that influences the way movement is perceived. To test the idea that the human body is special, Shiffrar and coworkers showed that when objects such as boards are used as stimuli, the likelihood of perceiving movement along the longer path doesn’t increase at lower rates of alternation, as it does for pictures of humans (Chatterjee, Freyd, & Shiffrar, 1996).

What is happening in the cortex when observers view apparent motion generated by pictures like the one in Figure 8.23? To find out, Jennifer Stevens and coworkers (2000) measured brain activation using the PET scan technique. They found that both movement through the head and movement around the head activated areas in the parietal cortex associated with movement. However, when the observers saw movement as occurring around the head, the motor cortex was activated as well. Thus, the motor cortex is activated when the perceived movements are humanly possible but isn’t activated when the perceived movements are not possible. This connection between the brain area associated with perceiving movement and the motor area reflects the close connection between perception and taking action that we discussed in Chapter 7.

Motion of Point-Light Walkers

Another approach to studying motion of the human body involves stimuli called **point-light walkers** that are created by placing small lights on people’s joints and then filming the patterns created by these lights when people walk and carry out other actions in the dark (Johansson, 1973, 1975) (**Figure 8.24**).

Perceptual Organization At the beginning of the chapter, we showed how movement can cause individual elements to become perceptually organized (see the camouflaged bird demonstration, page 177). Similarly, motion creates organization for point-light walkers. When the person wearing



Figure 8.24 A point-light walker is created by placing lights on a person’s joints and having the person walk in the dark so only the lights can be seen. © Cengage Learning

the lights is stationary, the lights look like a meaningless pattern. However, as soon as the person starts walking, with arms and legs swinging back and forth and feet moving in flattened arcs, first one leaving the ground and touching down, and then the other, the motion of the lights is immediately perceived as being caused by a walking person. This self-produced motion of a person or other living organism is called **biological motion**. **VL**

One reason we are particularly good at perceptually organizing the complex motion of an array of moving dots into the perception of a walking person is that we see biological

motion all the time. Every time you see a person walking, running, or behaving in any way that involves movement, you are seeing biological motion.

Brain Mechanisms Our ability to easily organize biological motions into meaningful perceptions led some researchers to suspect that there may be an area in the brain that responds to biological motion, just as there are areas such as the extrastriate body area (EBA) and fusiform face area (FFA) that are specialized to respond to bodies and faces, respectively.

Emily Grossman and Randolph Blake (2001) provided evidence supporting the idea of a specialized area in the brain for biological motion by measuring observers' brain activity as they viewed the moving dots created by a point-light walker (**Figure 8.25a**) and as they viewed dots that moved similarly to the point-light walker dots, but were scrambled so they did not result in the impression of a person walking (**Figure 8.25b**). They found that a small area in the superior temporal sulcus (STS) was more active when viewing biological motion than viewing scrambled motion in all eight of their observers. In another experiment, Grossman and Blake (2002) showed that other regions, such as the FFA, were activated more by biological motion than by scrambled motion, but that activity in the EBA did not distinguish between biological and scrambled motion. Based on these results, they concluded that there is a network of areas, which includes the STS and FFA, that is specialized for the perception of biological motion (also see Pelphrey et al., 2003).

One of the principles we have discussed in this book is that just showing that a structure responds to a specific type

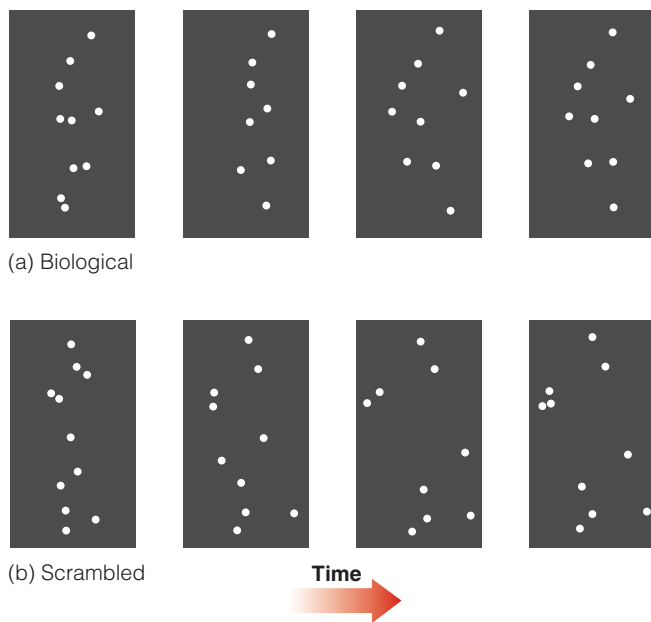


Figure 8.25 Frames from the stimuli used by Grossman and Blake (2001). (a) Sequence from the point-light walker stimulus. (b) Sequence from the scrambled point-light stimulus. From Grossman, E. D., & Blake, R. (2001). Brain activity evoked by inverted and imagined biological motion. *Vision Research*, 41, 1475–1482. With permission from Elsevier.

of stimulus does not prove that the structure is involved in *perceiving* that stimulus. Earlier in the chapter we described how Newsome used a number of different methods to show that the MT cortex is specialized for the perception of motion. In addition to showing that the MT cortex is *activated* by motion, he also showed that *perception* of motion is decreased by lesioning the MT cortex and is influenced by stimulating neurons in the MT cortex. Directly linking brain processes and perception enabled Newsome to conclude that the MT cortex is important for the perception of motion.

Just as Newsome showed that disrupting operation of the MT cortex decreases a monkey's ability to perceive the direction of moving dots, Emily Grossman and coworkers (2005) showed that disrupting operation of the STS in humans decreases the ability to perceive biological motion. Grossman accomplished this using a procedure called *transcranial magnetic stimulation*.

METHOD

Transcranial Magnetic Stimulation (TMS)

One way to investigate whether an area of the brain is involved in determining a particular function is to remove that part of the brain, as Newsome did in his studies of the MT cortex in monkeys. Of course, we cannot purposely remove a portion of a person's brain, but it is possible to temporarily disrupt the functioning of a particular area by applying a pulsating magnetic field using a stimulating coil placed over the person's skull (**Figure 8.26**). A series of pulses presented to a particular area of the brain for a few seconds interferes with brain functioning in that area for seconds or minutes. If a particular behavior is disrupted by the pulses, researchers conclude that the disrupted area of the brain is involved in that behavior.



Figure 8.26 TMS coil positioned to present a magnetic field to the back of the person's head.




Figure 8.27 (a) Biological motion stimulus. (b) Scrambled stimulus. (c) Biological motion stimulus with noise added. The dots corresponding to the walker are indicated by lines (which were not seen by the observer). (d) How the stimulus appears to the observer.

From Grossman, E. D., Battelli, L., & Pascual-Leone, A. (2005). Repetitive TMS over posterior STS disrupts perception of biological motion. *Vision Research*, 45, 2847–2853. With permission from Elsevier.

The observers in Grossman’s (2005) experiment viewed point-light stimuli for activities such as walking, kicking, and throwing (**Figure 8.27a**), and they also viewed scrambled point-light displays (**Figure 8.27b**). Their task was to determine whether a display was biological motion or

scrambled motion. This is normally an extremely easy task, but Grossman made it more difficult by adding extra dots to create “noise” (**Figure 8.27c** and **d**). The amount of noise was adjusted for each observer so that they could distinguish between biological and scrambled motion with 71 percent accuracy.

The key result of this experiment was that presenting transcranial magnetic stimulation to the area of the STS that is activated by biological motion caused a significant decrease in the observers’ ability to perceive biological motion. Such magnetic stimulation of other motion-sensitive areas, such as the MT cortex, had no effect on the perception of biological motion. From this result, Grossman concluded that normal functioning of the “biological motion” area, STS, is necessary for perceiving biological motion. This conclusion is also supported by studies showing that people who have suffered damage to this area have trouble perceiving biological motion (Battelli et al., 2003). What all of this means is that biological motion is more than just “motion”; it is a special type of motion that is served by specialized areas of the brain.

Representational Momentum: Motion Responses to Still Pictures

Look at the picture in **Figure 8.28**. Most people perceive this picture as a “freeze frame” of an action—skiing—that involves motion. It is not hard to imagine the person moving to a different location immediately after this picture was taken.



Alex Fetzler/CORBIS

Figure 8.28 A picture that creates implied motion.

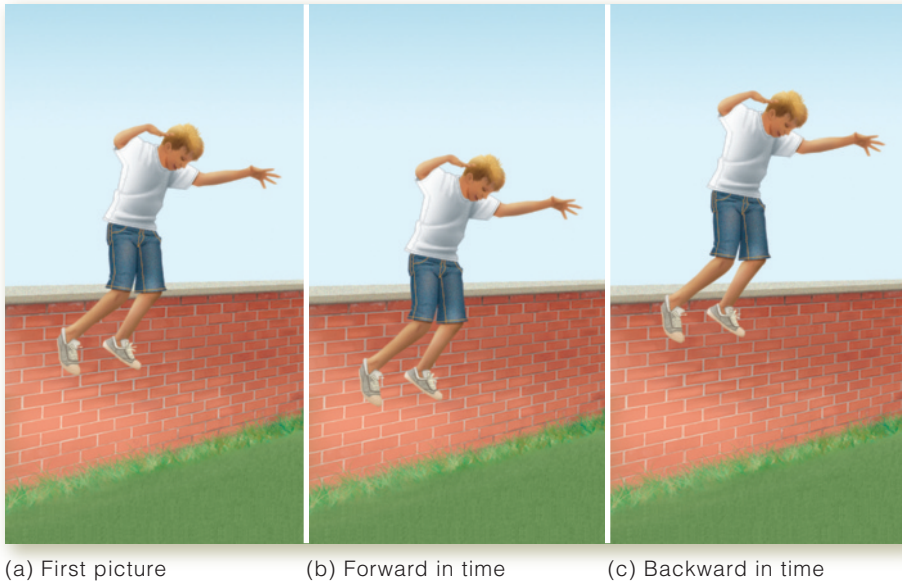


Figure 8.29 Stimuli like those used by Freyd (1983). See text for details. © Cengage Learning 2014

A situation such as this, in which a still picture depicts an action involving motion, is called **implied motion**.

Jennifer Freyd (1983) did an experiment involving implied motion by briefly showing observers pictures that depicted a situation involving motion, such as a person jumping off a low wall (**Figure 8.29a**). Freyd predicted that subjects looking at this picture would “unfreeze” the implied motion depicted in the picture and anticipate the motion that was about to happen. If this occurred, observers might “remember” the picture as depicting a situation that occurred slightly later in time. For the picture of the person jumping off the wall, that would mean the observers might remember the person as being closer to the ground (as in **Figure 8.29b**) than he was in the initial picture.

To test this idea, Freyd showed subjects a picture of a person in midair, like **Figure 8.29a**, and then after a pause, she showed her observers either (1) the same picture; (2) a picture slightly forward in time (the person who had jumped off the wall was closer to the ground, as in **Figure 8.29b**); or (3) a picture slightly backward in time (the person was farther from the ground, as in **Figure 8.29c**). The observers’ task was to indicate, as quickly as possible, whether the second picture was the same as or different from the first picture.

When Freyd compared the time it took for subjects to decide if the “time-forward” and “time-backward” pictures were different from the first picture they had seen, she found that subjects took longer to decide if the time-forward picture was the same or different. She concluded from this that the time-forward judgment was more difficult because her subjects had anticipated the downward motion that was about to happen and so confused the time-forward picture with what they had actually seen.

The idea that the motion depicted in a picture tends to continue in the observer’s mind is called **representational momentum** (David & Senior, 2000; Freyd, 1983). Representational momentum is an example of experience

influencing perception because it depends on our knowledge of the way situations involving motion typically unfold.

If implied motion causes an object to continue moving in a person’s mind, then it would seem reasonable that this continued motion might be reflected by activity in the brain. When Zoe Kourtzi and Nancy Kanwisher (2000) measured the fMRI response in the MT and MST cortex to pictures like the ones in **Figure 8.30**, they found that the area of the brain that responds to actual motion also responds to *pictures* of motion, and that implied-motion (IM) pictures caused a greater response than no-implied-motion (no-IM) pictures, at rest (R) pictures, or house (H) pictures. Thus, activity




Figure 8.30 Examples of pictures used by Kourtzi and Kanwisher (2000) to depict implied motion (IM), no implied motion (no-IM), at rest (R), and a house (H). The height of the bars below each picture indicates the average fMRI response of the MT cortex to that type of picture. From Kourtzi, Z., & Kanwisher, N., Activation in human MT/MST by static images with implied motion, *Journal of Cognitive Neuroscience*, 12, 1, January 2000, 48–55. © 2000 by Massachusetts Institute of Technology. All rights reserved. Reproduced by permission.

occurs in the brain that corresponds to the continued motion that implied-motion pictures create in a person's mind (also see Lorteije et al., 2006; Senior et al., 2000).

Building on the idea that the brain responds to implied motion, Jonathan Winawer and coworkers (2008) wondered whether still pictures that implied motion, like the one in Figure 8.28, would elicit a motion aftereffect (MAE). To test this, they did a psychophysical experiment in which they asked whether viewing still pictures showing implied motion in a particular direction can cause a motion aftereffect (MAE) in the opposite direction. We described one type of motion aftereffect at the beginning of the chapter by noting that after viewing the downward movement of a waterfall, nearby stationary objects appear to move upward. There is evidence that this occurs because prolonged viewing of the waterfall's downward motion decreases the activity of neurons that respond to downward motion, so more upward-motion neuronal activity remains (Barlow & Hill, 1963; Mather et al., 1998).

To determine whether implied motion stimuli would have the same effect, Winawer had his subjects observe a series of pictures showing implied motion. For a particular trial, subjects saw either a series of pictures that all showed movement to the right or a series of pictures that all showed movement to the left. After adapting to this series of pictures for 60 seconds, the subjects' task was to indicate the direction of movement of arrays of moving dots like the ones we described earlier (see Figure 8.16).

The key result of this experiment was that before observing the implied-motion stimuli, subjects were equally likely to perceive dot stimuli with zero coherence (all the dots moving in random directions) as moving to the left or to the right. However, after viewing photographs showing rightward implied motion, subjects were more likely to see the dots as moving to the left. After viewing leftward-implied motion, subjects were more likely to see the dots as moving to the right. Because this is the same result that would occur for adapting to real movement to the left or right, Winawer concluded that viewing implied motion in pictures decreases the activity of neurons selective to that direction of motion.

SOMETHING TO CONSIDER:

Event Perception

When you look out at a scene, you don't see an abstract arrangement of light, dark, and color. You see individual objects arranged relative to each other in space. This is the result of perceptual organization and perceptual segmentation, which we described in Chapter 5. When I see the papers, a coffee cup, keys, and a pen on the surface of the table in front of me, I am perceptually segregating this tabletop scene into separated objects.

But what does this have to do with perceiving movement? I'll answer that question by describing what I see when I look up from the array of objects on the table. A person enters the

coffee shop where I'm sitting, stops at the counter, has a brief conversation with the coffee barista behind the counter, who leaves and returns with coffee in a paper cup. The customer pushes down on the lid to make sure it is secure, pays for the coffee, drops a tip into the tip jar, turns around, and walks out the door. This short description, which represents only a small fraction of what is happening in the coffee shop, is a sequence of events unfolding in time. Just as we can segment a static scene into individual objects, we can segment ongoing behavior into a sequence of events, where an **event** is defined as a segment of time at a particular location that is perceived by observers to have a beginning and an ending (Zacks & Tversky, 2001; Zacks et al., 2009). An **event boundary** is the point in time when one event ends and another begins.

In our coffee shop scenario, placing an order with the coffee barista is an event; reaching out to accept the cup of coffee is an event; dropping change in the tip jar is an event; and so on. Our everyday life is a cascade of events, which can include our own behavior as well as our observations of the behaviors of others. The connection of events to motion perception becomes obvious when we consider that events almost always involve motion, and that changes in the nature of motion are often associated with event boundaries. One pattern of motion occurs when placing the order, another when reaching out for the coffee cup, and so on.

Jeffrey Zacks and coworkers (2009) have measured the connection between events and motion perception by having subjects watch films of common activities such as paying bills or washing dishes, and asking them to press a button when they believe one unit of meaningful activity ends and another begins (Newtson & Engquist, 1976; Zacks et al., 2001). When Zacks compared event boundaries to the actor's body movements measured with a motion tracking system, he found that event boundaries were more likely to occur when there was a change in the speed or acceleration of the actor's hands. From the results of this and other experiments, Zacks concluded that the perception of movement plays an important role in separating activities into meaningful events.

This brings us back to our example at the beginning of the chapter, in which we described the motions of a salesperson in a clothing store and noted that the person's motions indicated not only what she was doing (rearranging clothes) but also indicated when a new task began (helping a customer). Events, which are often defined by motion, follow one after the other to create our understanding of what is happening. **VL**

TEST YOURSELF 8.2

1. What is the evidence that the MT cortex is specialized for processing movement? Describe the series of experiments that used moving dots as stimuli and (a) recorded from neurons in the MT cortex, (b) lesioned the MT cortex, and (c) stimulated neurons in the MT cortex. What do the results of these experiments enable us to conclude about the role of the MT cortex in motion perception?

2. Describe the aperture problem—why the response of individual directionally selective neurons does not provide sufficient information to indicate the direction of motion. Also describe two ways that the brain might solve the aperture problem.
3. What is biological motion, and how has it been studied using point-light displays?
4. Describe experiments on apparent motion of a person's arm. How do the results differ for slow and fast presentations of the stimuli? How is the brain activated by slow and fast presentations?
5. Describe the experiments that have shown that an area in the STS is specialized for perceiving biological motion.
6. What is implied motion? Representational momentum? Describe behavioral evidence demonstrating representational momentum, physiological experiments that investigated how the brain responds to implied motion stimuli, and the experiment that used photographs to generate a motion aftereffect.
7. What is an event? What is the evidence that motion helps determine the location of event boundaries? What is the relation between events and our ability to predict what is going to happen next?

THINK ABOUT IT

1. We perceive real motion when we see things that are physically moving, such as cars on the road and people on the sidewalk. But we also see motion on TV, in movies, on our computer screens, and in electronic displays such as those in Las Vegas or Times Square. How are images presented in these situations in order to result in the perception of motion? (This may require some research.)
2. In the present chapter we have described a number of principles that also hold for object perception (Chapter 5). Find examples from Chapter 5 of the following (page numbers are for this chapter).
 - There are neurons that are specialized to respond to specific stimuli (182).
 - There are parallels between physiology and perception (187).
 - More complex stimuli are processed in higher areas of the cortex (192).
 - Experience can affect perception (190, 194).
3. Stark and Bridgeman explained the perception of movement that occurs when pushing gently on the eyelid by a corollary discharge signal generated when muscles are pushing back to counteract the push on the side of the eye. What if the push on the eyelid causes the eye to move, and the person sees the scene move? How would perception of the scene's movement in this situation be explained by corollary discharge theory? (p. 185)
4. We described how the representational momentum effect shows how knowledge can affect perception. Why could we also say that representational momentum illustrates an interaction between perception and memory? (p. 194)

KEY TERMS

Akinetopsia (p. 177)	Event (p. 195)	Motion aftereffect (p. 179)
Aperture problem (p. 189)	Event boundary (p. 195)	Motor signal (MS) (p. 183)
Apparent motion (p. 178)	Global optic flow (p. 182)	Optic array (p. 182)
Attentional capture (p. 177)	Illusory motion (p. 178)	Point-light walker (p. 191)
Biological motion (p. 191)	Image displacement signal (IDS) (p. 183)	Real motion (p. 178)
Coherence (p. 186)	Implied motion (p. 194)	Real-motion neuron (p. 186)
Comparator (p. 183)	Induced motion (p. 179)	Reichardt detector (p. 182)
Corollary discharge signal (CDS) (p. 183)	Local disturbance in the optic array (p. 182)	Representational momentum (p. 194)
Corollary discharge theory (p. 183)	Microstimulation (p. 187)	Shortest path constraint (p. 190)
		Waterfall illusion (p. 179)

CourseMate

Go to CengageBrain.com to access Psychology CourseMate, where you will find the Virtual Labs plus an interactive eBook, flashcards, quizzes, videos, and more.

Virtual Labs

The Virtual Labs are designed to help you get the most out of this course. The Virtual Lab icons direct you to specific media demonstrations and experiments designed to help you visualize what you are reading about. The numbers below indicate the number of the Virtual Lab you can access through Psychology CourseMate.

8.1 Motion Providing Organization: The Hidden Bird (p. 177)

How movement can cause an image to stand out from a complex background. (Courtesy of Michael Bach and David Regan)

8.2 Perceptual Organization: The Dalmatian (p. 178)

How a black and white pattern can be perceived as a Dalmatian. (Courtesy of Michael Bach)

8.3 Shape From Movement (p. 178)

How movement can create shape in an array of dots.

8.4 Larsen Experiment (p. 180)

Shows stimulus presentation for Larsen and colleagues' (2006) experiment. (Courtesy of Axel Larsen)

8.5 Corollary Discharge Model (p. 184)

A demonstration of how components of the model affect firing.

8.6 Motion Binding (p. 190)

Illustrates how adding an object to a display of four moving lines can influence how we perceive the motion of the lines. (Courtesy of Michael Bach)

8.7 Motion Perception in Depth (p. 190)

Narrated animation describing how conflicting right and left eye information creates perception of motion in depth. (Courtesy of Alex Huk)

8.8 Apparent Movement of the Human Body (p. 190)

Demonstration of possible and impossible apparent motion. (Courtesy of Maggie Shiffrar)

8.9 Biological Motion 1 (p. 191)

Illustrates how biological motion stimuli for a human walker change when gender, weight, and mood are varied. (Courtesy of Nikolaus Troje)

8.10 Biological Motion 2 (p. 191)

Illustrates biological motion stimuli for humans, cats, and pigeons and what happens when these stimuli are inverted, scrambled, and masked. (Courtesy of Nikolaus Troje)

8.11 Event Perception: Paying Bills (p. 195)

Shows person paying bills. Record below shows event segmentation and movement. (Courtesy of Jeffrey Zacks)

8.12 Event Perception: Working in Kitchen (p. 195)

Film of person working in kitchen. Record below shows coarse and fine event segmentation. (Courtesy of Jeffrey Zacks)



Perceiving Depth and Size

CHAPTER CONTENTS

Oculomotor Cues

Monocular Cues

Pictorial Cues
Motion-Produced Cues

Binocular Depth Information

Seeing Depth With Two Eyes
Binocular Disparity
Disparity (Geometrical) Creates Stereopsis (Perceptual)
The Correspondence Problem

The Physiology of Binocular Depth Perception

Perceiving Size

The Holway and Boring Experiment
Size Constancy

Visual Illusions

The Müller-Lyer Illusion
The Ponzo Illusion
The Ames Room
The Moon Illusion

SOMETHING TO CONSIDER: Depth Information Across Species

DEVELOPMENTAL DIMENSION: Infant Depth Perception

Using Binocular Disparity
Depth From Familiar Size
Depth From Cast Shadows

Think About It

◀ Our perception of depth is created by many sources of information in the environment. In this picture, the perception of depth is created by perspective convergence—the way parallel lines come together in the distance—and by a texture gradient—the way the brown and white structural elements of the building become more closely spaced farther in the distance. In this chapter we will consider many other sources of depth information and also the perception of size.

VL The Virtual Lab icons direct you to specific animations and videos designed to help you visualize what you are reading about. Virtual Labs are listed at the end of the chapter, keyed to the page on which they appear, and can be accessed through Psychology CourseMate.

Some Questions We Will Consider:

- How can we see far into the distance based on the flat image on the retina? (p. 231)
- Why do we see depth better with two eyes than with one eye? (p. 236)
- Why don't people appear to shrink in size when they walk away? (p. 248)

You can easily tell that this book is about 12 to 18 inches away and, when you look up at the scene around you, that other objects are located at distances ranging from your nose (very close!) to across the room, down the street, or even as far as the horizon, depending on where you are. What's amazing about this ability to see the distances of objects in your environment is that your perception of these objects, and the scene as a whole, is based on the flat two-dimensional image on your retina.

We can begin to appreciate the problem of perceiving depth based on two-dimensional information on the retina by considering two points on the scene in **Figure 10.1a**. Light is reflected from point T on the tree and from point H on the house onto points T and H on the retina at the back of the eye. Looking just at these points on the flat surface of the retina (**Figure 10.1b**), we have no way of knowing how far the light has traveled to reach each point. For all we know, the light stimulating either point on the retina could have come from 1 foot away or from a distant star. Clearly, we need to expand our view beyond single points on the retina to determine where objects are located in space.

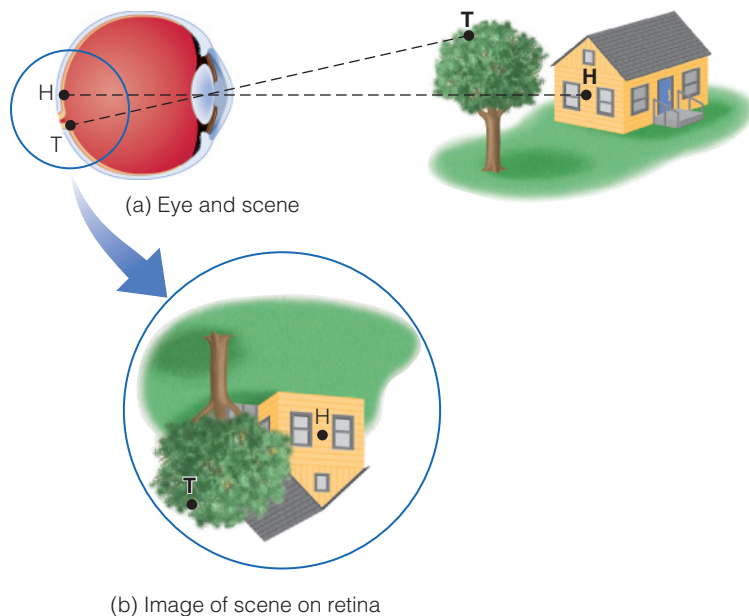


Figure 10.1 (a) In the scene, the house is farther away than the tree, but images of points H on the house and T on the tree both fall on the two-dimensional surface of the retina on the back of the eye. (b) These two points on the retinal image, considered by themselves, do not tell us the distances of the house and the tree.

© Cengage Learning

When we expand our view from two isolated points to the entire retinal image, we increase the amount of information available to us because now we can see the images of the house and the tree. However, because this image is two-dimensional, we still need to explain how we get from the flat image on the retina to the three-dimensional perception of the scene.

One way researchers have approached this problem is by the **cue approach to depth perception**, which focuses on identifying information in the retinal image that is correlated with depth in the scene. For example, when one object partially covers another object, as the tree in the foreground in Figure 10.1a covers part of the house, the object that is partially covered must be farther than the object that is covering it. This situation, which is called **occlusion**, is a cue that one object is in front of another. According to cue theory, we learn the connection between this cue and depth through our previous experience with the environment. After this learning has occurred, the association between particular cues and depth becomes automatic, and when these depth cues are present, we experience the world in three dimensions. A number of different types of cues that signal depth in a scene have been identified. We can divide these cues into three major groups:

1. *Oculomotor*. Cues based on our ability to sense the position of our eyes and the tension in our eye muscles.
2. *Monocular*. Cues that work with one eye.
3. *Binocular*. Cues that depend on two eyes.

Oculomotor Cues

The **oculomotor cues** are created by (1) convergence, the inward movement of the eyes that occurs when we look at nearby objects, and (2) accommodation, the change in the shape of the lens that occurs when we focus on objects at

various distances. The idea behind these cues is that we can *feel* the inward movement of the eyes that occurs when the eyes converge to look at nearby objects, and we feel the tightening of eye muscles that change the shape of the lens to focus on a nearby object. You can experience the feelings in your eyes associated with convergence and accommodation by doing the following demonstration. **VL**

DEMONSTRATION Feelings in Your Eyes

Look at your finger as you hold it at arm's length. Then, as you slowly move your finger toward your nose, notice how you feel your eyes looking inward and become aware of the increasing tension inside your eyes.

The feelings you experience as you move your finger closer are caused by (1) the change in convergence angle as your eye muscles cause your eyes to look inward, as in **Figure 10.2a**, and (2) the change in the shape of the lens as the eye accommodates to focus on a near object (Figure 2.4). If you move your finger farther away, the lens flattens, and your eyes move away from the nose until they are both looking straight ahead, as in **Figure 10.2b**. Convergence and accommodation indicate when an object is close and are useful up to a distance of about arm's length, with convergence being the more effective of the two (Cutting & Vishton, 1995; Mon-Williams & Tresilian, 1999; Tresilian et al., 1999).

Monocular Cues

Monocular cues work with only one eye. They include *accommodation*, which we have described under oculomotor cues; *pictorial cues*, which are sources of depth information in a

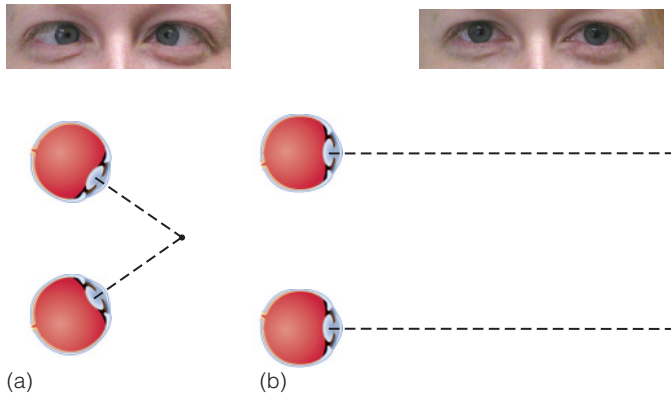


Figure 10.2 (a) Convergence of the eyes occurs when a person looks at something that is very close. (b) The eyes look straight ahead when the person observes something that is far away. © Cengage Learning 2014

two-dimensional picture; and *movement-based cues*, which are sources of depth information created by movement.

Pictorial Cues

Pictorial cues are sources of depth information that can be depicted in a picture, such as the illustrations in this book or an image on the retina (Goldstein, 2001b).

Occlusion We have already described the depth cue of occlusion. Occlusion occurs when one object hides or partially hides another from view. The partially hidden object is seen as being farther away, so the mountains in **Figure 10.3** are perceived as being farther away than the cactus and the

hill. Note that occlusion does not provide information about an object's distance. It indicates that the object that is partially covered is farther away than another object, but from occlusion alone we can't tell how much farther.

Relative Height In the photograph of the scene in **Figure 10.3a**, some objects are near the bottom of the frame and others nearer the top. The height in the frame of the photo corresponds to the height in our field of view, and objects that are higher in the field of view are usually farther away. This is illustrated in **Figure 10.3b**, in which dashed lines 1, 2, and 3 have been added under the front motorcycle, the rear motorcycle, and one of the telephone poles. Notice that dashed lines higher in the picture are under objects that are farther away. You can demonstrate this by looking out at a scene and placing your finger at the places where objects contact the ground. When you do this, you will notice that your finger is higher for farther objects. According to the cue of **relative height**, objects with their bases closer to the horizon are usually seen as being more distant. This means that being *higher* in the field of view causes objects on the *ground* to appear farther away (see lines 1, 2, and 3 in Figure 10.3b), whereas being *lower* in the field of view causes objects in the *sky* to appear farther away (see lines 4 and 5).

Relative Size According to the cue of **relative size**, when two objects are of equal size, the one that is farther away will take up less of your field of view than the one that is closer. This cue depends, to some extent, on a person's knowledge of physical sizes—for example, that the two telephone poles in Figure 10.3 are about the same size, as are the two motorcycles.



Figure 10.3 (a) A scene in Tucson, Arizona, containing a number of depth cues: occlusion (the cactus on the right occludes the hill, which occludes the mountain); relative height (the far motorcycle is higher in the field of view than the closer motorcycle); relative size (the far motorcycle and telephone pole are smaller than the near ones); and perspective convergence (the sides of the road converge in the distance). (b) 1, 2, and 3 indicate the increasing height in the field of view of the bases of the motorcycles and the far telephone pole, which reveals that being higher in the field of view causes objects on the ground to appear farther away; 4 and 5 reveal that being *lower* in the field of view causes objects in the sky to appear farther away.



Perugino, Pietro/ The Art Gallery Collection/Alamy

Figure 10.4 Pietro Perugino. *Christ Handing the Keys to St. Peter* (Sistine Chapel). The convergence of lines on the plaza illustrates perspective convergence. The sizes of the people in the foreground and middle ground illustrate relative size.

Perspective Convergence When you look down parallel railroad tracks that appear to converge in the distance, you are experiencing **perspective convergence**. This cue was often used by Renaissance artists to add to the impression of depth in their paintings, as in Pietro Perugino's painting in **Figure 10.4**. Notice that in addition to the perspective convergence provided by the lines on the plaza, Perugino has included people in the middle ground, enhancing the perception of depth further through the cue of relative size. Figure 10.3 illustrates both perspective convergence (the road) and relative size (the motorcycles) in our Tucson mountain scene.

Familiar Size We use the cue of **familiar size** when we judge distance based on our prior knowledge of the sizes of objects. We can apply this idea to the coins in **Figure 10.5**. If you are influenced by your knowledge of the actual size of dimes, quarters, and half-dollars, you would probably say that the dime is closer than the quarter. An experiment by William Epstein (1965) shows that under certain conditions, our knowledge of an object's size influences our perception of that object's distance. The stimuli in Epstein's experiment were equal-sized photographs of a dime, a quarter, and a



Figure 10.5 Drawings of the stimuli used in Epstein's (1965) familiar-size experiment. The actual stimuli were photographs that were all the same size as a real quarter. © Cengage Learning

half-dollar, which were positioned the same distance from an observer. By placing these photographs in a darkened room, illuminating them with a spot of light, and having subjects view them with one eye, Epstein created the illusion that these pictures were real coins.

When the observers judged the distance of each of the coin photographs, they estimated that the dime was closest, the quarter was farther than the dime, and the half-dollar was the farthest of all. The observers' judgments were influenced by their knowledge of the sizes of real dimes (small), quarters (larger), and half-dollars (largest). This result does not occur, however, when observers view the scene with both eyes, because, as we will see when we discuss binocular (two-eyed) vision, the use of two eyes provides information indicating the coins are at the same distance. The cue of familiar size is therefore most effective when other information about depth is absent (see also Coltheart, 1970; Schiffman, 1967).

Atmospheric Perspective **Atmospheric perspective** occurs when distant objects appear less sharp than nearer objects and often have a slight blue tint. **Figure 10.6** illustrates atmospheric perspective. The details in the foreground are sharp and well defined, but details become less and less visible as we look farther into the distance.

The farther away an object is, the more air and particles (dust, water droplets, airborne pollution) we have to look through, making objects that are farther away look less sharp and bluer than close objects.

The reason that farther objects look bluer is related to the reason the sky appears blue. Sunlight contains a distribution of all of the wavelengths in the spectrum, but the atmosphere preferentially scatters short-wavelength light (see Figure 9.4), which appears blue. This scattered light gives the sky its blue tint and also creates a veil of scattered light between us and objects we are looking at, although the blueness becomes obvious only



Bruce Goldstein

Figure 10.6 A scene on the coast of Maine showing the effect of atmospheric perspective.

when we are looking through a large distance or when there are more particles in the atmosphere to scatter the light.

If, instead of viewing this cliff along the coast of Maine, you were standing on the moon, where there is no atmosphere and hence no atmospheric perspective, far craters would not look blue and would look just as clear as near ones. But on Earth, there is atmospheric perspective, with the exact amount depending on the nature of the atmosphere.

Texture Gradient Another source of depth information is the **texture gradient**: Elements that are equally spaced in a scene appear to be more closely packed as distance increases, as in the scenes in **Figure 10.7**. Whether the closer packing



Londonistills.com/Alamy Limited

(a)




© Image Source/Corbis

(b)


Figure 10.7 Texture gradients created by marathon runners and flowers. The increasing fineness of texture as distance increases enhances the perception of depth.

occurs for marathon runners, flowers, or any other repeating elements seen in depth, the increasing fineness of texture as distance increases enhances the perception of depth.

Shadows Shadows—decreases in light intensity caused by the blockage of light—can provide information regarding the locations of these objects. Consider, for example, **Figure 10.8a**, which shows seven spheres and a checkerboard. In this picture, the location of the spheres relative to the checkerboard is unclear. They could be resting on the surface of the checkerboard or floating above it. But adding shadows, as shown in **Figure 10.8b**, makes the spheres' locations clearer—the ones on the left are resting on the checkerboard, and the



(a)



(b)

Figure 10.8 (a) Where are the spheres located in relation to the checkerboard? (b) Adding shadows makes their location clearer.

Courtesy of Pascal Mamassian.



(a)



(b)

Bruce Goldstein

Figure 10.9 (a) Early morning shadows emphasize the mountain's contours. (b) When the sun is overhead, the shadows vanish, and it becomes more difficult to see the mountain's contours.

ones on the right are floating above it. This illustrates how shadows can help determine the location of objects (Mamassian et al., 1998).

Shadows also enhance the three-dimensionality of objects. For example, shadows make the circles in Figure 10.8 appear spherical and help define some of the contours in the mountains in **Figure 10.9**, which appear three-dimensional in the early morning when there are shadows (**Figure 10.9a**), but flat in the middle of the day when the sun is directly overhead and there are no shadows (**Figure 10.9b**). **AVL**

Motion-Produced Cues

All of the cues we have described so far work if the observer is stationary. But once we start moving, new cues emerge that further enhance our perception of depth. We will describe two motion-produced cues: (1) motion parallax and (2) deletion and accretion.

Motion Parallax **Motion parallax** occurs when, as we move, nearby objects appear to glide rapidly past us, but more distant objects appear to move more slowly. Thus, when you look out the side window of a moving car or train, nearby objects appear to speed by in a blur, whereas objects that are farther away may appear to be moving only slightly.¹ We can understand why motion parallax occurs by noting how the image of a near object (the tree in **Figure 10.10a**) and a far object (the house in **Figure 10.10b**) move across the retina as an eye moves from position 1 to position 2 without rotating. First let's consider the tree: Figure 10.10a shows

¹If, when looking out the window, you keep your eyes fixed on one object, objects farther and closer than the object you are looking at appear to move in opposite directions.

one eye that moves from 1 to 2, so the tree's image moves all the way across the retina from T_1 to T_2 , as indicated by the dashed arrow. Figure 10.8b shows that the house's image moves a shorter distance, from H_1 to H_2 . Because the image of the tree travels a larger distance across the retina than the house, in the same amount of time, it appears to move more rapidly.

Motion parallax is one of the most important sources of depth information for many animals. The information provided by motion parallax has been used to enable human-designed mechanical robots to determine how far they are from obstacles as they navigate through the environment (Srinivasan & Venkatesh, 1997). Motion parallax is also widely used to create an impression of depth in cartoons and video games.

Deletion and Accretion As an observer moves sideways, some things become covered, and others become uncovered. Try the following demonstration.

DEMONSTRATION

Deletion and Accretion

Close one eye. Position your hands as shown in **Figure 10.11**, so your right hand is at arm's length and your left hand at about half that distance, just to the left of the right hand. Then as you look at your right hand, move your head sideways to the left, being sure to keep your hands still. As you move your head, your left hand appears to cover your right hand. This covering of the farther right hand is **deletion**. If you then move your head back to the right, the nearer hand moves back and uncovers the right hand. This uncovering of the far hand is **accretion**. Deletion and accretion occur all the time as we move through the environment and create information that the object or surface being covered and uncovered is farther away (Kaplan, 1969).

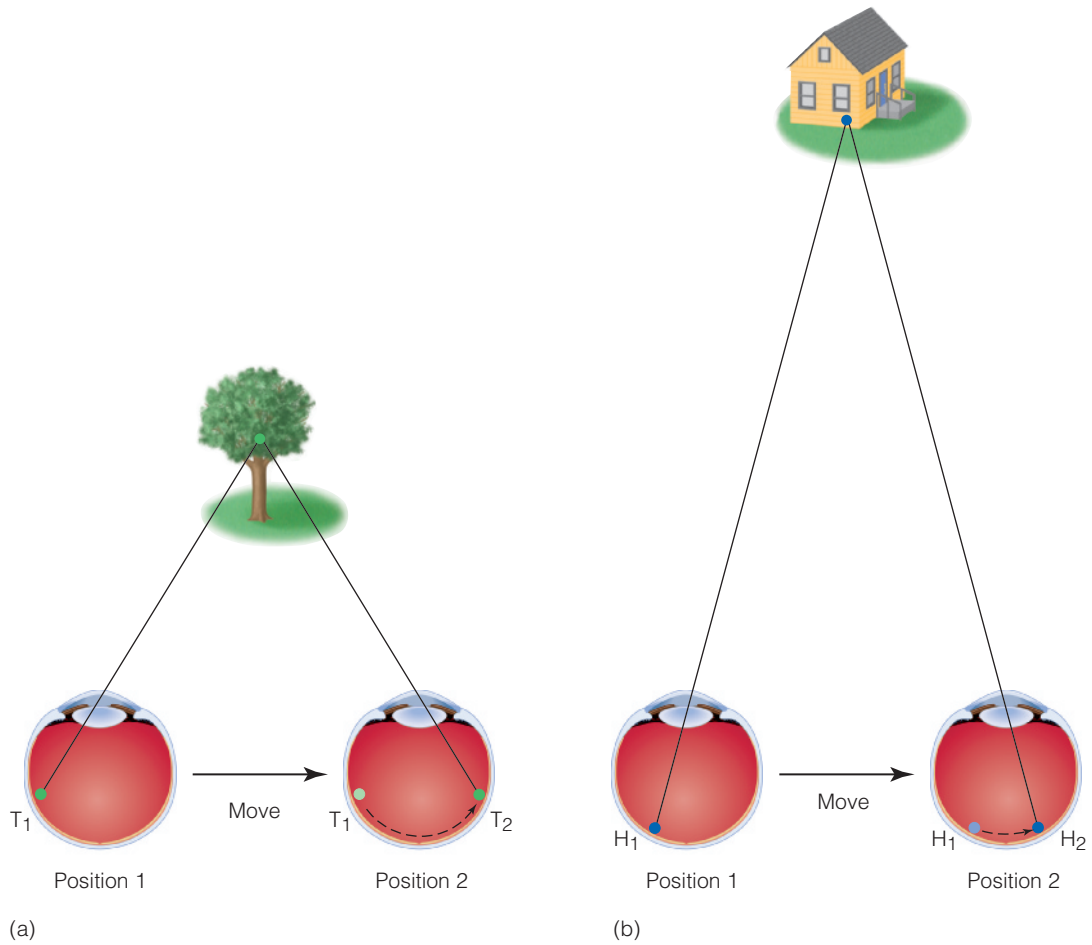


Figure 10.10 One eye moving past (a) a nearby tree; (b) a far-away house. Because the tree is closer, its image moves farther across the retina than the image of the house. © Cengage Learning



Figure 10.11 Position of the hands for “Deletion and Accretion” demonstration.

Our discussion so far has described a number of the cues that contribute to our perception of depth. As shown in **Table 10.1**, these cues work over different distances: some only at close range (accommodation, convergence); some at close and medium ranges (motion parallax, deletion and accretion);

TABLE 10.1 Range of Effectiveness of Different Depth Cues

DEPTH INFORMATION	0-2 METERS	2-20 METERS	ABOVE 30 METERS
Deletion and accretion	✓	✓	
Occlusion	✓	✓	✓
Relative size	✓	✓	✓
Accommodation and convergence	✓		
Motion parallax	✓	✓	
Relative height		✓	✓
Atmospheric perspective			✓

Source: Based on Cutting & Vishton, 1995.

some at long range (atmospheric perspective, relative height); and some at the whole range of depth perception (occlusion, relative size; Cutting & Vishton, 1995).

Binocular Depth Information

One of the myths I heard sometime during my childhood was that you need both eyes to perceive depth. I soon figured out that this wasn't true, because when I closed one eye, I could still tell what was near and what was far away. But sometimes myths can be partially true. Although it is possible to use monocular cues to see depth, there is something qualitatively different about the depth perception experienced when using both eyes. Two-eyed depth perception, called **stereoscopic vision**, involves mechanisms that take into account differences in the images formed on the left and right eyes. The following demonstration illustrates these differences.

DEMONSTRATION

Two Eyes: Two Viewpoints

Close your right eye. Hold a finger on your left hand at arm's length. Position a right-hand finger about a foot away, so it covers the other finger. Then open the right eye and close the left. When you switch eyes, how does the position of your front finger change relative to the rear finger?

When you switched from looking with your left eye to your right, you probably noticed that the front finger appeared to move to the left relative to the far finger. **Figure 10.12** diagrams what happened on your retinas. The green line in **Figure 10.12a** shows that when the left eye is open, the images of the near and far fingers were lined up with the same place on the retina. This occurred because

you were looking directly at both objects, so both images would fall on the foveas of the left eye. The green lines in **Figure 10.12b** show that when the right eye was open, the image of the far finger still fell on the fovea because you were looking at it, but the image of the near finger was now off to the side.

Whereas the fingers were lined up relative to the left eye, the right eye "looks around" the near finger, so the far finger becomes visible. These different viewpoints for the two eyes is the basis of **stereoscopic depth perception**—depth perception created by input from both eyes. Before describing these mechanisms, we will consider what it means to say that stereoscopic depth perception is qualitatively different from monocular depth perception.

Seeing Depth With Two Eyes

Three-dimensional movies were introduced to the public on a large scale in the 1950s, when audiences were introduced to 3-D glasses, and *The House of Wax* became the highest grossing 3-D movie. Three-dimensional movies soon lost their allure, both because of the quality of the stories and the inconvenience of wearing 3-D glasses, and were relegated mainly to short features shown in theme parks. But recently, with the development of better 3-D technology and films like *Avatar* (2009) and *Hugo* (2011), 3-D movies have become a standard fixture of moviegoing, with 3-D TV sets not far behind (more on this later). If you have seen a 3-D movie, it is easy to appreciate the added dimension provided by stereoscopic depth. Scenes seen in 3-D appear to have added depth compared to 2-D, with objects sometimes appearing to jut far out from the screen.

The main reason for the difference between our perception of 2-D and 3-D movies is shown in **Figure 10.13**. Even though we view a 2-D movie with both eyes, the screen




Figure 10.12 Location of images on the retina for the “Two Eyes: Two Viewpoints” demonstration. See text for explanation. © Cengage Learning




Figure 10.13 (a) When we view a two-dimensional movie, the left and right eyes receive essentially the same images, so depth is indicated only by monocular pictorial depth cues. (b) When viewing a 3-D movie, the left and right eyes receive different images, so stereoscopic depth perception occurs.

© Cengage Learning 2014

is flat, so both eyes receive essentially the same images (**Figure 10.13a**). Thus any depth perceived in these movies results from monocular or pictorial depth cues.

The situation for 3-D movies is different, because 3-D technology causes the left and right eyes to receive slightly different images, as shown in **Figure 10.13b**. These different views duplicate what happens in the real 3-D world, which we see from two viewing positions, as illustrated in the finger demonstration.

Another way to appreciate the qualitative difference between monocular depth perception and stereoscopic depth perception is to consider the story of Susan Barry, a neuroscientist at Mt. Holyoke College. Her story—first described by neurologist Oliver Sacks, who dubbed her “Stereo Sue” (Sachs, 2006, 2010), and then in her own book, *Fixing My Gaze* (Barry, 2011)—begins with Susan’s childhood eye problems. She was cross-eyed, so when she looked at something with one eye, the other eye would be looking somewhere else. For most people, both eyes aim at the same place and work in coordination with each other, but in Susan’s case, the input was uncoordinated. Situations such as this, along with a condition called “walleye” in which the eyes look out, are forms of **strabismus**, or misalignment of the eyes. When this occurs, the visual system suppresses vision in one of the eyes to avoid double vision, so the person sees the world with only one eye at a time.

Susan had a number of operations as a child, which made it more difficult to detect her strabismus, but her vision was

still dominated by one eye. Although her perception of depth was achieved through monocular cues, she was able to get along quite well. She could drive, play softball, and do most of the things people with stereoscopic vision can do. For example, she describes her vision in a college classroom as follows:

I looked around. The classroom didn’t seem entirely flat to me. I knew that the student sitting in front of me was located between me and the blackboard because the student blocked my view of the blackboard. When I looked outside the classroom window, I knew which trees were located further away because they looked smaller than the closer ones. (Barry, 2011, Chapter 1)

Although Susan could use the monocular cues she describes above to perceive depth, her knowledge of the neuroscience literature and various other experiences she describes in her book led her to realize that she was still seeing with one eye despite her childhood operations. She therefore consulted an optometrist, who confirmed her one-eyed vision and assigned eye exercises designed to improve the coordination between her two eyes. These exercises enabled Susan to coordinate her eyes, and one day after leaving the optometrist’s office, she had her first experience with stereoscopic depth perception, which she describes as follows:

I got into my car, sat down in the driver’s seat, placed the key in the ignition, and glanced at the steering

wheel. It was an ordinary steering wheel against an ordinary dashboard, but it took on a whole new dimension that day. The steering wheel was floating in its own space, with a palpable volume of empty space between the wheel and the dashboard. I closed one eye and the steering wheel looked “normal” again; that is, it lay flat just in front of the dashboard. I reopened the closed eye, and the steering wheel floated before me. (Barry, 2011, Chapter 6)

From that point on, Susan had many more experiences that astounded her, much as someone who had never experienced stereoscopic vision might react if they could put on 3-D movie glasses and suddenly begin seeing in stereoscopic three dimensions. It is important to note that Susan didn’t suddenly gain stereovision equivalent to that experienced by a person with stereoscopic vision from birth. Her stereovision occurred first for nearby objects and then, as her training progressed, was extended to farther distances. But what she did experience dramatically illustrates the richness that stereoscopic vision adds to the experience of depth perception.

Binocular Disparity

Binocular disparity, the differences in the images on the left and right retinas, is the basis of the stereoscopic vision Susan experienced. We now look more closely at the information

on the left and right retinas that the brain uses to create an impression of depth.

Corresponding Retinal Points We begin by introducing **corresponding retinal points**—points on the retina that overlap if the eyes are superimposed on each other (**Figure 10.14**). We can illustrate corresponding points by considering the observer in **Figure 10.15a**, who is looking directly at Julie. **Figure 10.15b** shows where Julie’s images are located on the observer’s retinas. Because the observer is looking directly




Figure 10.14 Corresponding points on the two retinas. To determine corresponding points, imagine that the left eye is slid on top of the right eye. F indicates the fovea, where the image of an object occurs when an observer looks directly at the object, and A is a point in the peripheral retina. Images on the fovea always fall on corresponding points. Notice that the A’s, which also fall on corresponding points, are the same distance from the fovea in the left and right eyes. © Cengage Learning




Figure 10.15 (a) An observer looking at Julie’s face, with a tree off to the side. (b) The observer’s eyes, showing where the images of Julie and the tree fall on each eye. Julie’s images fall on the fovea, so they are on corresponding points. The arrows indicate that the tree’s images are located the same distances from the fovea in the two eyes, so they are also on corresponding points. The dashed blue line is the horoapter. The images of objects that are on the horoapter fall on corresponding points. © Cengage Learning 2014

at Julie, her images fall on the observer's foveas on both eyes, indicated by the green dots. The two foveas are corresponding points, so Julie's images fall on corresponding points.

In addition, the images of other objects also fall on corresponding points. Consider, for example, the tree in Figure 10.15b. The tree's images are on the same place relative to the foveas—to the left and at the same distance (indicated by the arrows). This means that the tree's images are on corresponding points. (If you were to slide the eyes on top of each other, Julie's images would overlap, and the tree's images would overlap.) Thus, whatever a person is looking at directly (like Julie) falls on corresponding points, and some other objects (like the tree) fall on corresponding points as well. Julie, the tree, and any other objects that fall on corresponding points are located on a surface called the **horopter**. The blue dashed lines in Figure 10.15a and 10.15b show part of the horopter.

Absolute Disparity Indicates Distance From the Horopter

The images of objects that are *not* on the horopter fall on **noncorresponding points**. The degree to which these objects *deviate* from falling on corresponding points is called **absolute disparity**. This is illustrated in **Figure 10.16a**, which shows Julie again, with her images on corresponding points, and a new character, Bill, whose images are on noncorresponding points. The amount of absolute disparity, which is called the **angle of disparity**, is indicated by the red arrow,

which shows the angle between the corresponding point for the left-eye image of Bill (red dot) and where the image is actually located (black dot).

Figure 10.16b shows that the angle of disparity is greater for objects at greater distances from the horopter. The observer is still looking at Julie, and Bill is where he was before, but now we have added Dave, who is located even farther from the horopter than Bill. When we compare Dave's angle of disparity in this figure (blue arrow) to Bill's in Figure 10.16a (red arrow), we see that Dave's disparity is greater. (The same thing also happens for objects farther away than the horopter, with greater distance also associated with greater disparity.) The angle of disparity therefore provides information about an object's distance from the horopter, with greater angles of disparity indicating greater distances from the horopter.

Relative Disparity Is Related to Objects' Positions Relative to Each Other

Let's now consider what happens when the observer shifts his gaze from one object to another. When the observer is looking at Julie (**Figure 10.17a**), Julie's images fall on the observer's foveas (so Julie's disparity is zero), but the images of Bill fall on noncorresponding points (so there is disparity). But when the observer shifts his gaze to Bill (**Figure 10.17b**), Bill's images fall on the foveas (so Bill's disparity is now zero) and Julie's images fall on noncorresponding points (so there is disparity).

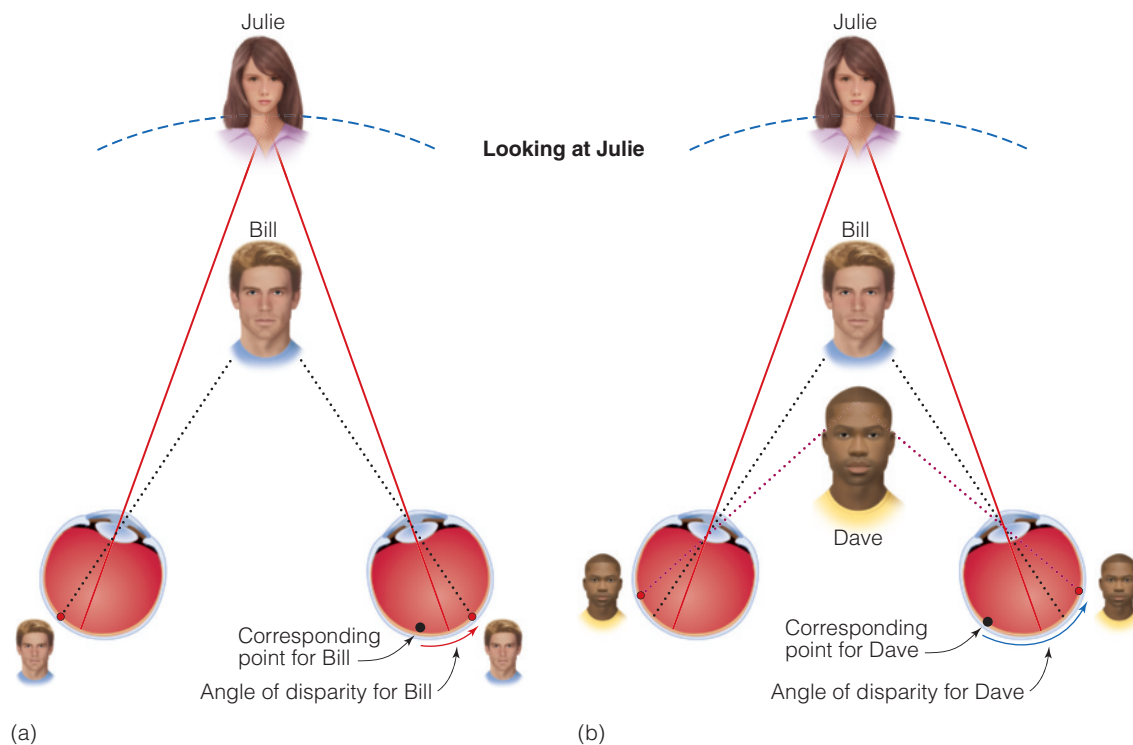


Figure 10.16 (a) When the observer looks at Julie, Julie's images fall on corresponding points. Bill's images fall on noncorresponding points. The angle of disparity, indicated by the red arrow, is determined by measuring the angle between where the corresponding point for Bill's image would be located (black dot) and where Bill's image is actually located (red dot). (b) Dave has been added to Figure 10.16a. Dave's angle of disparity (blue arrow) is greater than Bill's, because Dave is located farther from the horopter. © Cengage Learning 2014




Figure 10.17 Absolute disparities change when an observer's gaze shifts from one place to another. (a) When the observer looks at Julie, the disparity of her images is zero. Bill's angle of disparity is indicated by the arrow. (b) When the observer looks at Bill, the disparity of Bill's images becomes zero. Julie's angle of disparity is indicated by the arrow. Because one of the disparities in each pair is zero, the arrows indicate the *difference in disparity* between Julie's and Bill's images. Note that the difference in disparity is the same in (a) and (b). This means that the *relative disparity* of Julie and Bill remains the same as the observer looks at different places. © Cengage Learning 2014

If we compare the two situations in 10.17a and b, we notice that the difference in absolute disparities between Julie and Bill (indicated by the lengths of the arrows) is the same in both situations. The *difference* in absolute disparities of objects in a scene, called **relative disparity**, remains the same as an observer looks around a scene. Relative disparity helps indicate where objects in a scene are located relative to one another. As we will see below, there is evidence that both absolute and relative disparity information is represented by neural activity in the visual system.

Disparity (Geometrical) Creates Stereopsis (Perceptual)

We have seen that both absolute and relative disparity information contained in the images on the retinas provides information indicating an object's distance from where the observer is looking. Notice, however, that our description of disparity has focused on *geometry*—looking at where objects' images fall on the retina—but has not mentioned *perception*, the observer's experience of an object's depth or its relation to other objects in the environment (**Figure 10.18**).




Figure 10.18 Disparity is related to geometry—the locations of images on the retina. *Stereopsis* is related to perception—the experience of depth created by disparity. © Cengage Learning 2014

We consider the relationship between disparity and what observers perceive by introducing **stereopsis**—the impression of depth that results from information provided by binocular disparity. An example of stereopsis is provided by the depth effect achieved by the **stereoscope**, a device introduced by the



(a) Left eye image



(b) Right eye image

Bruce Goldstein

Figure 10.19 The two images of a stereoscopic photograph. The difference between the two images, such as the distances between the front cactus and the window in the two views, creates retinal disparity. This creates a perception of depth when the left image is viewed by the left eye and the right image is viewed by the right eye.

physicist Charles Wheatstone (1802–1875), which produces a convincing illusion of depth by using two slightly different pictures. This device, extremely popular in the 1800s and reintroduced as the View Master in the 1940s, presents two photographs made with a camera with two lenses separated by the same distance as the eyes. The result is two slightly different views, like those shown in **Figure 10.19**. The stereoscope presents the left picture to the left eye and the right picture to the right eye. This creates the same binocular disparity that occurs when a person views the scene naturally, so that slightly different images appear in the left and right eyes.

The principle behind the stereoscope is also used in 3-D movies. The left-eye and right-eye images are presented superimposed on the screen, slightly displaced from one another, to create disparity. There are a number of ways of achieving this. One way is to color one image red and the other green and to view the film through glasses with a red filter for one eye and a green filter for the other eye. Another way of separating the left and right images, which has been

used in movies like *Avatar* and *Hugo*, is to create the left and right images from polarized light—light waves that vibrate in only one orientation. One image is polarized so its vibration is vertical, and the other is polarized so its vibration is horizontal. Viewing the film through polarizing lenses, which let vertically polarized light into one eye and horizontally polarized light into the other eye, creates the disparity that results in three-dimensional perception. This method creates better color than the red–green method, which results in little or no variation in color.

Similar techniques are used to create 3-D perception of TV images, but with some variations based on the way TV images are created. The main methods are illustrated in **Figure 10.20**. The **passive method** works according to the principles we have described for 3-D movies, with two superimposed polarized images viewed through polarizing glasses (**Figure 10.20a**). The **active method** alternates the left-eye and right-eye images on the screen 30 or more times a second. This method is called active because the viewing glasses




Figure 10.20 Three types of 3-D TV. See text for details. © Cengage Learning 2014

have a shutter system that is synchronized with the alternation of images occurring on the TV screen, so the shutter for the left eye opens when the left-eye image is present on the screen, and the shutter for the right eye opens when the right-eye image is present (Figure 10.20b). A disadvantage of this method is that the glasses are expensive, and some people report headaches after extended viewing; an advantage is that better image quality may be possible than with passive viewing.

In a third method, called **lenticular projection**, the screen is coated with a film that contains two sets of lenses that direct different images to the left and right eyes (Figure 10.20c). You may have experienced lenticular images in postcards that show different images when viewed at different angles or that create a 3-D impression when viewed from one viewing point. An advantage of this method is that glasses are not required. Disadvantages are that the effect works best when viewed from a particular viewpoint, and viewing the images causes motion sickness in some people. The technology of 3-D television is developing so rapidly that some of the disadvantages described here may have become less important by the time you are reading this book!

Returning to our consideration of whether disparity creates stereopsis, we can point out that all of the methods we have been describing use disparity to create 3-D perception. However, these examples don't conclusively prove that disparity creates stereopsis, because images such as those

in Figure 10.19 also contain potential depth cues, such as occlusion and relative height, which could contribute to our perception of depth. In order to show that disparity alone can result in depth perception, Bela Julesz (1971) created a stimulus called the *random-dot stereogram*, which contains no pictorial cues.

By creating stereoscopic images of random-dot patterns, Julesz showed that observers can perceive depth in displays that contain no depth information other than disparity. Two such random-dot patterns, which together constitute a **random-dot stereogram**, are shown in Figure 10.21. These patterns were constructed by first generating two identical random-dot patterns on a computer and then shifting a square-shaped section of the dots one or more units to the side.

In the stereogram in Figure 10.21a, a section of dots from the pattern on the left has been shifted one unit to the right to form the pattern on the right. This shift is too subtle to be seen in the dot patterns, but we can understand how it is accomplished by looking at the diagrams below the dot patterns (Figure 10.21b). In these diagrams, the black dots are indicated by 0's, A's, and X's and the white dots by 1's, B's, and Y's. The A's and B's indicate the square-shaped section where the shift is made in the pattern. Notice that the A's and B's are shifted one unit to the right in the right-hand pattern. The X's and Y's indicate areas uncovered by the shift that must be filled in with new black dots and white dots to complete the pattern.



(a)

1	0	1	0	1	0	0	1	0	1
1	0	0	1	0	1	0	1	0	0
0	0	1	1	0	1	1	0	1	0
0	1	0	A	A	B	B	1	0	1
1	1	1	B	A	B	A	0	0	1
0	0	1	A	A	B	A	0	1	0
1	1	1	B	B	A	B	1	0	1
1	0	0	1	1	0	1	1	0	1
1	1	0	0	1	1	0	1	1	1
0	1	0	0	0	1	1	1	1	0

1	0	1	0	1	0	0	1	0	1
1	0	0	1	0	1	0	1	0	0
0	0	1	1	0	1	1	0	1	0
0	1	0	Y	A	A	B	B	0	1
1	1	1	X	B	A	B	A	0	1
0	0	1	X	A	A	B	A	1	0
1	1	1	Y	B	B	A	B	0	1
1	0	0	1	1	0	1	1	0	1
1	1	0	0	1	1	0	1	1	1
0	1	0	0	0	1	1	1	1	0

(b)

Figure 10.21 (a) A random-dot stereogram. (b) The principle for constructing the stereogram. See text for an explanation. © Cengage Learning

Even though it is not possible to tell that the dots have been shifted when looking at Figure 10.21a, the visual system detects a difference when the left image is presented to the left eye and the right image to the right eye. The disparity created by the shifted section results in perception of a small square floating above the background. Because binocular disparity is the only depth information present in these stereograms, disparity alone must be causing the perception of depth.

Psychophysical experiments, particularly those using Julesz's random-dot stereograms, show that retinal disparity creates a perception of depth. But before we can fully understand the mechanisms responsible for depth perception, we must answer one more question: How does the visual system match the parts of the images in the left and right eyes that correspond to one another? This is called the **correspondence problem**, and as we will see, it has still not been fully explained.

The Correspondence Problem

Let's return to the stereoscopic images of Figure 10.19. When we view this image in a stereoscope, we see different parts of the image at different depths because of the disparity between images on the left and right retinas. Thus, the cactus and the window appear to be at different distances when viewed through the stereoscope because they create different amounts of disparity. But in order for the visual system to calculate this disparity, it must compare the images of the cactus on the left and right retinas and the images of the window on the left and right retinas. This is the correspondence problem. How does the visual system match up the images in the two eyes?

A possible answer to this question is that the visual system may match the images on the left and right retinas on the basis of the specific features of the objects. For example, the upper-left windowpane on the left could be matched with the upper-left pane on the right, and so on. Explained in this way, the solution seems simple: Most things in the world are quite discriminable from one another, so it is easy to match an image on the left retina with the image of the same thing on the right retina. But what about images in which matching similar points would be extremely difficult, as with Julesz's random-dot stereogram?

You can appreciate the problem involved in matching similar parts of a stereogram by trying to match up the points in the left and right images of the stereogram in Figure 10.21. Most people find this to be an extremely difficult task, involving switching their gaze back and forth between the two pictures and comparing small areas of the pictures one after another. But even though matching similar features on a random-dot stereogram is much more difficult and time-consuming than matching features in the real world, the visual system somehow matches similar parts of the two stereogram images, calculates their disparities, and creates a perception of depth. A number of proposals, all too complex to describe here, have been put forth to explain

how the visual system solves the correspondence problem, but a totally satisfactory answer has yet to be proposed (see Blake & Wilson, 1991; Menz & Freeman, 2003; Ohzawa, 1998; Ringbach, 2003).

The Physiology of Binocular Depth Perception

The idea that binocular disparity provides information for the positions of objects in space implies that there should be neurons that signal different amounts of disparity. Research beginning in the 1960s and 1970s revealed neurons that respond to absolute disparity (Barlow et al., 1967; Hubel & Wiesel, 1970). These neurons are called **binocular depth cells** or **disparity-selective cells**. A given cell responds best when stimuli presented to the left and right eyes create a specific amount of absolute disparity (Uka & DeAngelis, 2003). **Figure 10.22** shows a **disparity tuning curve** for one of these neurons. This particular neuron responds best when the left and right eyes are stimulated to create an absolute disparity of about 1 degree. Further research has shown that there are also neurons higher up in the visual system that respond to relative disparity (Parker, 2007).

Brain-imaging experiments on humans show that a number of different areas are activated by stimuli that create binocular disparity (Backus et al., 2001; Kwee et al., 1999; Ts'o et al., 2001). Experiments on monkeys have determined that neurons sensitive to absolute disparity are found in the primary visual receiving area, and neurons sensitive to relative disparity are found in the temporal lobe and other areas. Apparently, depth perception involves a number of stages of processing, beginning in the primary visual cortex and extending to many different areas in both the ventral and dorsal streams (Parker, 2007).

The relationship between binocular disparity and the firing of binocular depth cells is an example of the stimulus-physiology relationship in the diagram of the perceptual process in **Figure 10.23** (orange arrow). This diagram, which we introduced in Chapter 1 (see Figure 1.10) and repeated




Figure 10.22 Disparity tuning curve for a neuron sensitive to absolute disparity. This curve indicates the neural response that occurs when stimuli presented to the left and right eyes create different amounts of disparity. © Cengage Learning 2014

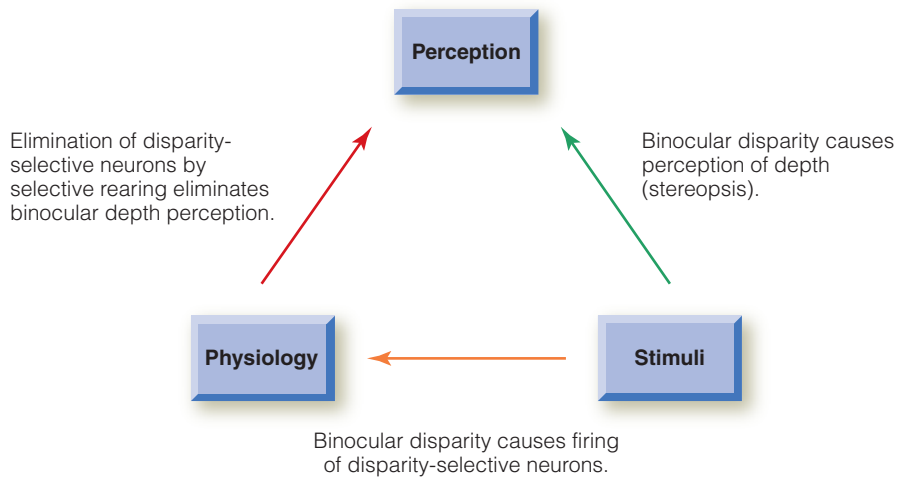


Figure 10.23 The three relationships in the perceptual process, as applied to binocular disparity. We have described experiments relating disparity to perception (green arrow) and relating disparity to physiological responding (orange arrow). The final step is to determine the relationship between physiological responses to disparity and perception (red arrow). This has been studied by selective rearing, which eliminates disparity-selective neurons, as well as by other methods described in the text. © Cengage Learning 2014

in Chapter 8 (see Figure 8.17), also depicts two other relationships. The stimulus–perception relationship (green arrow) is the relationship between binocular disparity and the perception of depth. The final relationship, between physiology and perception (red arrow), involves demonstrating a connection between disparity-selective neurons and depth perception. This has been achieved in a number of ways.

An early demonstration of a connection between binocular neurons and perception involved the selective rearing procedure we described in our discussion of the relationship between feature detectors and perception in Chapter 3 (see page 66). Applying this procedure to depth perception, Randolph Blake and Helmut Hirsch (1975) reared cats so that their vision was alternated between the left and right eyes every other day during the first 6 months of their lives. After this 6-month period of presenting stimuli to just one eye at a time, Blake and Hirsch recorded from neurons in the cat’s cortex and found that (1) these cats had few binocular neurons, and (2) they were not able to use binocular disparity to perceive depth. Thus, eliminating binocular neurons eliminates stereopsis and confirms what everyone suspected all along—that disparity-selective neurons are responsible for stereopsis (also see Olson & Freeman, 1980).

Another technique that has been used to demonstrate a link between neural activity and depth perception is microstimulation, a procedure in which a small electrode is inserted into the cortex and an electrical charge is passed through the electrode to activate the neurons near the electrode (Cohen & Newsome, 2004). (See Method: Microstimulation in Chapter 8, page 187) In Chapter 8 we described research that showed that stimulating neurons that respond best to specific directions of movement shifts a monkey’s perception of moving dots toward that direction of movement. Gregory DeAngelis and coworkers (1998) demonstrated the same effect for depth perception by training monkeys to indicate the depth created by presenting images with different absolute disparities. Presumably, the monkey perceived

depth because the disparate images on the monkey’s retinas activated disparity-selective neurons in the cortex. But what would happen if microstimulation were used to activate a different group of disparity-selective neurons?

Neurons that are sensitive to the same disparities tend to be organized in clusters, so stimulating one of these clusters activates a group of neurons that respond best to a specific disparity. When DeAngelis and coworkers stimulated neurons that were tuned to a disparity different from what was indicated by the images on the retina, the monkey shifted its depth judgment toward the disparity signaled by the stimulated neurons (**Figure 10.24**). The results of the selective rearing and the microstimulation experiments indicate that binocular depth cells are a physiological mechanism responsible for depth perception, thus providing the physiology–perception relationship of the perceptual process in Figure 10.23.




Figure 10.24 While the monkey was observing a random-dot stereogram, DeAngelis and coworkers (1998) stimulated neurons in the monkey’s cortex that were sensitive to a particular amount of disparity. This stimulation shifted the monkey’s perception of the depth of the field of dots from position 1 to position 2. © Cengage Learning

TEST YOURSELF 10.1

1. What is the basic problem of depth perception, and how does the cue approach deal with this problem?
2. What monocular cues provide information about depth in the environment?
3. What do comparing the experience of viewing 3-D and 2-D movies and the experiences of “Stereo Sue” tell us about what binocular vision adds to our perception of depth?
4. What is binocular disparity? What is the difference between absolute disparity and relative disparity? How are absolute and relative disparity related to the depths of objects in a scene?
5. What is stereopsis? What is the evidence that disparity creates stereopsis?
6. What does perception of depth from a random-dot stereogram demonstrate?
7. What is the correspondence problem? Has this problem been solved?
8. Describe each of the relationships in the perceptual process of Figure 10.23, and provide examples for each relationship that has been determined by psychophysical and physiological research on depth perception.

Perceiving Size

We discuss size perception in this chapter because our perception of size can be affected by our perception of depth. This link between size perception and depth perception is graphically illustrated by the example of whiteout, a treacherous weather condition faced by helicopter pilots flying across snow-covered terrain. The following description, based on an actual incident at an Antarctic research facility, illustrates the effect of whiteout on size perception:

As Frank pilots his helicopter across the Antarctic wastes, blinding light, reflected down from thick cloud cover above and up from the pure white blanket of snow below, makes it difficult to see the horizon, details on the surface of the snow, or even up from down. He is aware of the danger because he has known pilots dealing with similar conditions who flew at full power directly into the ice. He thinks he can make out a vehicle on the snow far below, and he drops a smoke grenade to check his altitude. To his horror, the grenade falls only three feet before hitting the ground. Realizing that what he thought was a truck was actually a small box, Frank pulls back on the controls and soars up, his face drenched in sweat, as he comprehends how close he just came to becoming another whiteout fatality.




Figure 10.25 When a helicopter pilot loses the ability to perceive distance in a “whiteout,” a small box that is close can be mistaken for a truck that is far away.

This account illustrates that our ability to perceive an object’s size can sometimes be drastically affected by our ability to perceive the object’s distance. A small box seen close up can, in the absence of accurate information about its distance, be misperceived as a large truck seen from far away (**Figure 10.25**). The idea that we can misperceive size when accurate depth information is not present was demonstrated in a classic experiment by A. H. Holway and Edwin Boring (1941).

The Holway and Boring Experiment

Observers in Holway and Boring’s experiment sat at the intersection of two hallways and saw a luminous *test circle* when looking down the right hallway and a luminous *comparison circle* when looking down the left hallway (**Figure 10.26**).




Figure 10.26 Setup of Holway and Boring’s (1941) experiment. The observer changes the diameter of the comparison circle in the left corridor to match his or her perception of the size of test circles presented in the right corridor. Each test circle has a visual angle of 1 degree and is presented separately. This diagram is not drawn to scale. The actual distance of the far test circle was 100 feet.

The comparison circle was always 10 feet from the observer, but the test circles were presented at distances ranging from 10 feet to 120 feet. An important property of the fixed-in-place comparison circle was that its size could be adjusted. The observer's task on each trial was to adjust the diameter of the comparison circle in the left corridor to match his or her perception of the sizes of the various test circles presented in the right corridor.

An important feature of the test stimuli in the right corridor was that they all cast exactly the same-sized image on the retina. We can understand how this was accomplished by introducing the concept of visual angle.

What Is Visual Angle? *Visual angle* is the angle of an object relative to the observer's eye. **Figure 10.27a** shows how we determine the visual angle of a stimulus (a person, in this example) by extending lines from the person to the lens of the observer's eye. The angle between the lines is the visual angle. Notice that the visual angle depends both on the size of the stimulus and on its distance from the observer, so when the person moves closer, as in **Figure 10.27b**, the visual angle becomes larger.

The visual angle tells us how large the object will be on the back of the eye. There are 360 degrees around the entire circumference of the eyeball, so an object with a visual angle of 1 degree would take up 1/360 of this circumference—about 0.3 mm in an average-sized adult eye. One way to get a feel for visual angle is to fully extend your arm and look at your thumb, as the woman in **Figure 10.28** is doing. The approximate visual angle of the *width* of the thumb at arm's length is 2 degrees. Thus, an object that is exactly covered by the thumb held at arm's length, such as the phone in Figure 10.28, has a visual angle of approximately 2 degrees.

This "thumb technique" provides a way to determine the approximate visual angle of any object in the environment. It also illustrates an important property of visual angle: A small object that is near (like the thumb) and a larger object that is far (like the phone) can have the same visual angle. An extreme example of this is illustrated in **Figure 10.29**, which shows a photograph taken by Jennifer, a student in my sensation and perception class. To take this picture, Jennifer adjusted the distance between her fingers so that the Eiffel Tower just fit between them. When she did this,

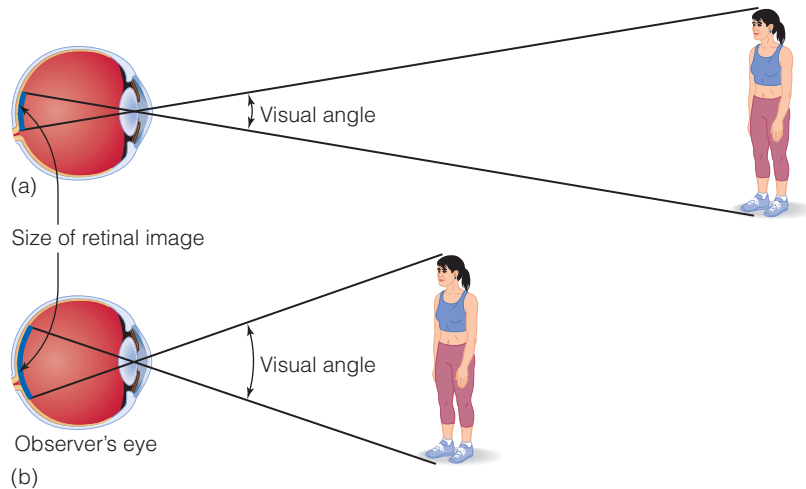


Figure 10.27 (a) The visual angle depends on the size of the stimulus (the woman in this example) and its distance from the observer. (b) When the woman moves closer to the observer, the visual angle and the size of the image on the retina increase. This example shows that halving the distance between the stimulus and observer doubles the size of the image on the retina. © Cengage Learning




Figure 10.28 The "thumb" method of determining the visual angle of an object. When the thumb is at arm's length, it has a visual angle of about 2 degrees. The woman's thumb covers the width of her phone, so the visual angle of the phone, from the woman's point of view, is 2 degrees. Note that the visual angle will change if the distance between the woman and the phone changes. © Cengage Learning



Figure 10.29 The visual angle between the two fingers is the same as the visual angle of the Eiffel tower.

the space between her fingers, which were about a foot away, had the same visual angle as the Eiffel Tower, which was miles away.

How Holway and Boring Tested Size Perception in a Hallway The idea that objects with different sizes can have the same visual angle was used in the creation of the test circles in Holway and Boring’s experiment. As shown in Figure 10.26, small circles that were positioned close to the observer and larger circles that were positioned farther away all had visual angles of 1 degree. Because objects with the same visual angle create the same-sized image on the retina, all of the test circles had the same-sized image on the observers’ retinas, no matter where in the hallway they were located.

In the first part of Holway and Boring’s experiment, many depth cues were available, including binocular disparity, motion parallax, and shading, so the observer could easily judge the distance of the test circles. The results, plotted in **Figure 10.30**, show that when the observers viewed a large test circle that was located far away (far circle in Figure 10.26), they made the comparison circle large (point F in Figure 10.30); when they viewed a small test circle that was located nearby (near circle in Figure 10.26), they made the comparison circle small (point N in Figure 10.30). Thus, when good depth cues were present, the observer’s judgments of the size of the circles matched the physical sizes of the circles.

Holway and Boring then determined how eliminating depth information would affect the observer’s judgments




Figure 10.30 Results of Holway and Boring’s experiment. The dashed line labeled “physical size” is the result that would be expected if the observers adjusted the diameter of the comparison circle to match the actual diameter of each test circle. The line labeled “visual angle” is the result that would be expected if the observers adjusted the diameter of the comparison circle to match the visual angle of each test circle. © Cengage Learning

of size. They did this by having the observer view the test circles with one eye, which eliminated binocular disparity (line 2 in Figure 10.30); then by having the observer view the test circles through a peephole, which eliminated motion parallax (line 3); and finally by adding drapes to the hallway to eliminate shadows and reflections (line 4). Each time some depth information was eliminated, the observer’s judgments of the sizes of the test circles became less accurate. When all depth information was eliminated, the observer’s perception of size was determined not by the actual size of the test circles but by the relative sizes of the circle’s images on the observer’s retinas.

Because all of the test circles in Holway and Boring’s experiment had the same retinal size, eliminating depth information caused them to be perceived as being about the same size. Thus, the results of this experiment indicate that size estimation is based on the actual sizes of objects when there is good depth information (blue lines), but that size estimation is strongly influenced by the object’s visual angle when depth information is eliminated (red lines).

An example of size perception that is determined by visual angle is our perception of the sizes of the sun and the moon, which, by cosmic coincidence, have the same visual angle. The fact that they have identical visual angles becomes most obvious during an eclipse of the sun. Although we can see the flaming corona of the sun surrounding the moon, as shown in **Figure 10.31**, the moon’s disk almost exactly covers the disk of the sun.

If we calculate the visual angles of the sun and the moon, the result is 0.5 degrees for both. As you can see in Figure 10.31, the moon is small (diameter 2,200 miles) but close (245,000 miles from Earth), whereas the sun is large (diameter 865,400 miles) but far away (93 million miles




Figure 10.31 The moon's disk almost exactly covers the sun during an eclipse because the sun and the moon have the same visual angles. © Cengage Learning

from Earth). Even though these two celestial bodies are vastly different in size, we perceive them to be the same size because, as we are unable to perceive their distance, we base our judgment on their visual angles.

In yet another example, we perceive objects viewed from a high-flying airplane as very small. Because we have no way of accurately estimating the distance from the airplane to the ground, we perceive size based on objects' visual angles, which are very small because we are so high up.

Size Constancy

One of the most obvious features of the scene in **Figure 10.32**, on the campus of the University of Arizona, is that looking down the row of palm trees, each more distant tree becomes smaller in the picture. If you were standing on campus observing this scene, the more distant trees would appear to take up



Figure 10.32 All of the palm trees appear to be the same size when viewed in the environment, even though the farther ones have a smaller visual angle.

less of your field of view, as in the picture, but at the same time you would not perceive the farther tree as shorter than the near trees. Even though the far trees take up less of your field of view (or to put it another way, have a smaller *visual angle*), they appear constant in size. The fact that our perception of an object's size is relatively constant even when we view the object from different distances is called **size constancy**.

To introduce the idea of size constancy to my perception class, I ask someone in the front row to estimate my height when I am standing about 3 feet away. Their guess is usually accurate, around 5 feet 9 inches. I then take one large step back so I am now twice as far away and ask the person to estimate my height again. It probably doesn't surprise you that the second estimate of my height is about the same as the first. The point of this demonstration is that even though my image on the person's retina becomes half as large when I double my distance (compare Figures 10.27a and b), I do not appear to shrink to less than 3 feet tall, but still appear to be my normal size. The following demonstration illustrates size constancy in another way.

DEMONSTRATION

Perceiving Size at a Distance

Hold a quarter between the fingertips of each hand so you can see the faces of both coins. Hold one coin about a foot from you and the other at arm's length. Observe the coins with both of your eyes open and note their sizes. Under these conditions, most people perceive the near and far coins as being approximately the same size. Now close one eye, and holding the coins so they appear side-by-side, notice how your perception of the size of the far coin changes so that it now appears smaller than the near coin. This demonstrates how size constancy is decreased under conditions of poor depth information.

Although students often propose that size constancy works because we are familiar with the sizes of objects, research has shown that observers can accurately estimate

the sizes of unfamiliar objects viewed at different distances (Haber & Levin, 2001).

Size Constancy as a Calculation The link between size constancy and depth perception has led to the proposal that size constancy is based on a mechanism called **size–distance scaling** that takes an object’s distance into account (Gregory, 1966). Size–distance scaling operates according to the equation $S = K(R \times D)$, where S is the object’s perceived size, K is a constant, R is the size of the retinal image, and D is the perceived distance of the object. (Since we are mainly interested in R and D , and K is a scaling factor that is always the same, we will omit K in the rest of our discussion).

According to the size–distance equation, as a person walks away from you, the size of the person’s image on your retina (R) gets smaller, but your perception of the person’s distance (D) gets larger. These two changes balance each other, and the net result is that you perceive the person’s size (S) as staying the same.

DEMONSTRATION

Size–Distance Scaling and Emmert’s Law

You can demonstrate size–distance scaling to yourself by looking back at Figure 8.12 in Chapter 8 (page 184). Look at the center of the circle for about 60 seconds. Then look at the white space to the side of the circle. If you blink, you should see the circle’s afterimage floating in front of the page. Before the afterimage fades, also look at a wall far across the room. You should see that the size of the afterimage depends on where you look. If you look at a distant surface, such as the far wall of the room, you see a large afterimage that appears to be far away. If you look at a near surface, such as the page of this book, you see a small afterimage that appears to be close.




Figure 10.33 illustrates the principle underlying the effect you just experienced, which was first described by Emmert in 1881. Staring at the circle bleached a small circular area of visual pigment on your retina. This bleached area of the retina determined the retinal size of the afterimage and remained constant no matter where you were looking.

The perceived size of the afterimage, as shown in Figure 10.33, is determined by the distance of the surface against which the afterimage is viewed. This relationship between the apparent distance of an afterimage and its perceived size is known as **Emmert’s law**: The farther away an afterimage appears, the larger it will seem. This result follows from our size–distance scaling equation, $S = R \times D$. The size of the bleached area of pigment on the retina (R) always stays the same, so that increasing the afterimage’s distance (D) increases the magnitude of $R \times D$. We therefore perceive the size of the afterimage (S) as larger when it is viewed against the far wall.

The size–distance scaling effect demonstrated by the afterimage demonstration is working constantly when we look at objects in the environment, with the visual system taking both an object’s size in the field of view (which determines retinal size) and its distance into account to determine our perception of its size. This process, which is happening constantly without any effort on our part, helps us perceive a stable environment. Just think of how confusing it would be if objects appeared to shrink or expand just because we happen to be viewing them from different distances. Luckily, because of size constancy, this doesn’t happen.

Other Information for Size Perception Although we have been stressing the link between size constancy and depth perception and how size–distance scaling works, other sources of information in the environment also help us achieve size

Figure 10.33 The principle behind the observation that the size of an afterimage increases as the afterimage is viewed against more distant surfaces. © Cengage Learning



Bruce Goldstein

Figure 10.34 The size of this wheel becomes apparent when it is compared to the person. If the wheel were seen in isolation, it would be difficult to know that it is so large.



Bruce Goldstein

Figure 10.35 Two cylinders resting on a texture gradient. The fact that the bases of both cylinders cover the same portion of a paving stone indicates that the two cylinders are the same size.

constancy. One source of information for size perception is relative size. We often use the sizes of familiar objects as a yardstick to judge the size of other objects, as in **Figure 10.34**, in which the size of the woman indicates that the wheel is very large. This idea that our perception of the sizes of objects can be influenced by the sizes of nearby objects explains why we often fail to appreciate how tall basketball players are, when all we see for comparison are other basketball players. But as soon as a person of average height stands next to one of these players, the player's true height becomes evident.

Another source of information for size perception is the relationship between objects and texture information on the ground. We saw that a texture gradient occurs when elements that are equally spaced in a scene appear to be more closely packed as distance increases (Figure 10.7). **Figure 10.35** shows two cylinders sitting on a texture gradient formed by a cobblestone road. Even if we have trouble perceiving the depth of the near and far cylinders, we can tell that they are the same size because their bases both cover the same portion of a paving stone.

Visual Illusions

Visual illusions fascinate people because they demonstrate how our visual system can be “tricked” into seeing inaccurately (Bach & Poloschek, 2006). We have already described a number of types of illusions. Illusions of lightness include Mach bands (page 56), in which small changes in lightness are seen near a border even though no changes are present in the physical pattern of light; simultaneous contrast (page 58)

and White's illusion (page 59), in which two physically identical fields can appear different; and the Hermann grid (page 56), in which small gray spots are seen that aren't there in the light. Attentional effects include change blindness (page 139), in which two alternating scenes appear similar even though there are differences between them. Illusions of motion are those in which stationary stimuli are perceived as moving (page 178).

We will now describe some illusions of size—situations that lead us to misperceive the size of an object. We will see that some explanations of these illusions involve the connection we have described between the perception of size and the perception of depth. We will also see that some of the most familiar illusions have yet to be fully explained. A good example of this situation is provided by the Müller-Lyer illusion.

The Müller-Lyer Illusion

In the **Müller-Lyer illusion**, the right vertical line in **Figure 10.36** appears to be longer than the left vertical line, even though they are both exactly the same length (measure them). A number of different explanations have been proposed to explain this illusion. An influential early explanation involves size–distance scaling.

Misapplied Size Constancy Scaling Why does the Müller-Lyer display cause a misperception of size? Richard Gregory (1966) explains the illusion on the basis of a mechanism he calls **misapplied size constancy scaling**. He points out that size constancy normally helps us maintain a stable perception of objects by taking distance into account




Figure 10.36 The Müller-Lyer illusion. Both lines are actually the same length. © Cengage Learning

(as expressed in the size–distance scaling equation). Thus, size constancy scaling causes a 6-foot-tall person to appear 6 feet tall no matter what his distance. Gregory proposes, however, that the very mechanisms that help us maintain stable perceptions in the three-dimensional world sometimes create illusions when applied to objects drawn on a two-dimensional surface.

We can see how misapplied size constancy scaling works by comparing the left and right lines in Figure 10.36 to the left and right lines that have been superimposed on the corners in **Figure 10.37**. Both lines are the same size, but according to Gregory the lines appear to be at different distances because the fins on the right line in Figure 10.37 make this line look like part of an inside corner of a room, and the fins on the left line make this line look like part of a corner viewed from outside. Because inside corners appear to “recede” and outside corners “jut out,” our size–distance scaling mechanism treats the inside corner as if it is farther away, so the term D in the equation $S = R \times D$ is larger and this line therefore appears longer. (Remember that the retinal sizes, R , of the two lines are the same, so perceived size, S , is determined by the perceived distance, D .)

At this point, you could say that although the Müller-Lyer figures may remind Gregory of inside and outside corners, they don’t look that way to you (or at least they didn’t until Gregory told you to see them that way). But according to Gregory, it is not necessary that you be consciously aware that these lines can represent three-dimensional structures; your perceptual system unconsciously takes the depth information contained in the Müller-Lyer figures into account, and your size–distance scaling mechanism adjusts the perceived sizes of the lines accordingly.

Gregory’s theory of visual illusions has not, however, gone unchallenged. For example, figures like the dumbbells in



Bruce Goldstein

Figure 10.37 According to Gregory (1966), the Müller-Lyer line on the left corresponds to an outside corner, and the line on the right corresponds to an inside corner. Note that the two vertical lines are the same length (measure them!).

Figure 10.38, which contain no obvious perspective or depth, still result in an illusion. And Patricia DeLucia and Julian Hochberg (1985, 1986, 1991; Hochberg, 1987) have shown that the Müller-Lyer illusion occurs for a three-dimensional display like the one in **Figure 10.39**, in which it is obvious that the spaces between the two sets of fins are not at different depths. (Measure distances x and y to convince yourself that they are the same.) You can experience this effect for yourself by doing the following demonstration.

DEMONSTRATION

The Müller-Lyer Illusion With Books

Pick three books that are the same size and arrange two of them with their corners making a 90-degree angle and standing in positions A and B, as shown in Figure 10.39. Then, without using a ruler, position the third book at position C, so that distance x appears to be equal to distance y . Check your placement, looking down at the books from the top and from other angles as well. When you are satisfied that distances x and y appear about equal, measure the distances with a ruler. How do they compare?

If you set distance y so that it was smaller than distance x , this is exactly the result you would expect from the two-dimensional Müller-Lyer illusion, in which the distance between the outward-facing fins appears enlarged compared to the distance between the inward-facing fins. You can also




Figure 10.38 The “dumbbell” version of the Müller-Lyer illusion. As in the original Müller-Lyer illusion, the two straight lines are actually the same length. © Cengage Learning




Figure 10.39 A three-dimensional Müller-Lyer illusion. The 2-foot-high wooden “fins” stand on the floor. Although the distances x and y are the same, distance y appears larger, just as in the two-dimensional Müller-Lyer illusion. © Cengage Learning

duplicate the illusion shown in Figure 10.39 with your books by using your ruler to make distances x and y equal. Then, notice how the distances actually appear. The fact that we can create the Müller-Lyer illusion by using three-dimensional stimuli such as these, along with demonstrations like the dumbbell in Figure 10.38, is difficult for Gregory’s theory to explain.

Conflicting Cues Theory R. H. Day (1989, 1990) has proposed the **conflicting cues theory**, which states that our perception of line length depends on two cues: (1) the actual length of the vertical lines, and (2) the overall length of the figure. According to Day, these two conflicting cues are integrated to form a compromise perception of length. Because the overall length of the right figure in Figure 10.36 is larger due to its outward-oriented fins, the vertical line appears larger.

Another version of the Müller-Lyer illusion, shown in **Figure 10.40**, results in the perception that the space between the dots is greater in the lower figure than in the upper figure, even though the distances are actually the same. According to Day’s conflicting cues theory, the space in the lower figure appears greater because the overall extent of the figure is greater. Notice that conflicting cues theory can also be applied to the dumbbell display in Figure 10.38. Thus, although Gregory believes that depth information is involved in determining illusions, Day rejects this idea and proposes that cues for length are what is important. Let’s now look at some more examples of illusions and the mechanisms that have been proposed to explain them.

The Ponzo Illusion

In the **Ponzo** (or railroad track) **illusion**, shown in **Figure 10.41**, both animals are the same size on the page, and so have the same visual angle, but the one on top appears longer. According to Gregory’s misapplied scaling explanation, the top animal appears bigger because of depth information provided by the converging railroad tracks that make




Figure 10.40 An alternate version of the Müller-Lyer illusion. We perceive that the distance between the dots in (a) is less than the distance in (b), even though the distances are the same. (From Day, 1989.) © Cengage Learning



Figure 10.41 The Ponzo (or railroad track) illusion. The two animals are the same length on the page (measure them), but the upper one appears larger. Courtesy of Mary Bravo

the top animal appear farther away. Thus, just as in the Müller-Lyer illusion, the scaling mechanism corrects for this apparently increased depth (even though there really isn't any, because the illusion is on a flat page), and we perceive the top animal to be larger. (Also see Prinzmetal et al., 2001; Shimamura & Prinzmetal, 1999, for another explanation of the Ponzo illusion.) **VL**

The Ames Room

The **Ames room** causes two people of equal size to appear very different in size (Ittelson, 1952). In **Figure 10.42**, you can see that the woman on the right looks much taller than the woman on the left. This perception occurs even though both women are actually about the same height. The reason for this erroneous perception of size lies in the construction of the room. The shapes of the wall and the windows at the rear of the room make it look like a normal rectangular room when viewed from a particular observation point; however, as shown in the diagram in **Figure 10.43**, the Ames room is actually shaped so that the left corner of the room is almost twice as far from the observer as the right corner.

What's happening in the Ames room? The construction of the room causes the woman on the left to have a much smaller visual angle than the one on the right. We think that we are looking into a normal rectangular room at two women who appear to be at the same distance, so we perceive the



Figure 10.42 The Ames room. Both women are actually the same height, but the woman on the right appears taller because of the distorted shape of the room. (The Exploratorium/S. Schwartzberg.)




Figure 10.43 The Ames room, showing its true shape. The person on the left is actually almost twice as far away from the observer as the person on the right; however, when the room is viewed through the peephole, this difference in distance is not seen. In order for the room to look normal when viewed through the peephole, it is necessary to enlarge the left side of the room. © Cengage Learning

one with the smaller visual angle as shorter. We can understand why this occurs by returning to our size–distance scaling equation, $S = R \times D$. Because the *perceived* distance (D) is the same for the two women, but the size of the retinal image (R) is smaller for the woman on the left, her perceived size (S) is smaller.

Another explanation for the Ames room is based not on size–distance scaling but on relative size. The relative size explanation states that our perception of the size of the two women is determined by how they fill the distance between the bottom and top of the room. Because the woman on the right fills the entire space and the woman on the left occupies only a little of it, we perceive the woman on the right as taller (Sedgwick, 2001).

The Moon Illusion

You may have noticed that when the moon is on the horizon, it appears much larger than when it is higher in the sky. This enlargement of the horizon moon compared to the elevated moon, shown in **Figure 10.44**, is called the **moon illusion**. When I discuss this in class, I first explain that visual angles of the horizon moon and elevated moon are the same. This must be so because the moon's physical size (2,200 miles in diameter) stays the same (obviously) and it remains the same distance from Earth (245,000 miles) throughout the night; therefore, the moon's visual angle must be constant. (If you are still skeptical, photograph the horizon moon and the elevated moon with a digital camera. When you compare the two images, you will find that the diameters in the resulting two pictures are identical. Or you can view the moon through a quarter-inch-diameter hole held at about arm's length. For most people, the moon just fits inside this hole, wherever it is in the sky.)



Figure 10.44 An artist's conception of how the moon is perceived when it is on the horizon and when it is high in the sky. Note that the visual angle of the horizon moon is depicted as larger than the visual angle of the moon high in the sky. This is because the picture is simulating the illusion. In the environment, the visual angles of the two moons are the same. © Cengage Learning

Once students are convinced that the moon's visual angle remains the same throughout the night, I ask why they think the moon appears larger on the horizon. One common response is "When the moon is on the horizon, it appears closer, and that is why it appears larger." When I ask why it appears closer, I often receive the explanation "Because it appears larger." But saying "It appears larger because it appears closer, and it appears closer because it appears larger" is clearly a case of circular reasoning that doesn't really explain the moon illusion.

One explanation that isn't circular is called the **apparent distance theory**. This theory does take distance into account, but in a way opposite to our hypothetical student's explanation. According to apparent distance theory, the moon on the horizon appears more distant because it is viewed across the filled space of the terrain, which contains depth information; but when the moon is higher in the sky, it appears less distant because it is viewed through empty space, which contains little depth information.

The idea that the horizon is perceived as farther away than the sky overhead is supported by the fact that when people estimate the distance to the horizon and the distance to the sky directly overhead, they report that the horizon appears to be farther away. That is, the heavens appear "flattened" (**Figure 10.45**).

The key to the moon illusion, according to apparent distance theory, is that the horizon moon and the elevated moon have the same visual angle, but because the horizon moon is seen against the horizon, which appears farther than the zenith sky, it appears larger. This follows from the size-distance scaling equation, $S = R \times D$. Retinal size, R , is the same for both locations of the moon (remember that the visual angle is always the same no matter where the moon appears in the sky), so the moon that appears farther away will appear larger. This is the principle we invoked in the Emmert's law demonstration to explain why an afterimage appears larger if it is viewed against a faraway surface.

Just as the near and far afterimages in the Emmert's law demonstration have the same visual angles, so do the horizon and elevated moons. The afterimage that appears on the far wall simulates the horizon moon; the circle appears farther away, so your size-distance scaling mechanism makes it appear larger. The afterimage that is viewed on a close surface simulates the elevated moon; the circle appears closer,




Figure 10.45 When observers are asked to consider the sky as a surface and to compare the distance to the horizon (H) and the distance to the top of the sky on a clear moonless night, they usually say that the horizon appears farther away. This results in the "flattened heavens" shown here. © Cengage Learning

so your scaling mechanism makes it appear smaller (King & Gruber, 1962).

Lloyd Kaufman and Irvin Rock (1962a, 1962b) have done a number of experiments that support the apparent distance theory. In one of their experiments, they showed that when the horizon moon was viewed over the terrain, which made it seem farther away, it appeared 1.3 times larger than the elevated moon; however, when the terrain was masked off so that the horizon moon was viewed through a hole in a sheet of cardboard, the illusion vanished (Kaufman & Rock, 1962a, 1962b; Rock & Kaufman, 1962).

Some researchers, however, are skeptical of the apparent distance theory. They question the idea that the horizon moon appears farther, as shown in the flattened heavens effect in Figure 10.45, because some observers see the horizon moon as floating in space in front of the sky (Plug & Ross, 1994).

Another theory of the moon illusion is the **angular size contrast theory**, which states that the moon appears smaller when it is surrounded by larger objects. Thus, when the moon is elevated, the large expanse of sky surrounding it makes it appear smaller. However, when the moon is on the horizon, less sky surrounds it, so it appears larger (Baird et al., 1990).

Even though scientists have been proposing theories to explain the moon illusion for hundreds of years, there is still no agreement on an explanation (Hershenson, 1989). Apparently a number of factors are involved, in addition to the ones we have considered here, including atmospheric perspective (looking through haze on the horizon can increase size perception), color (redness increases perceived size), and oculomotor factors (convergence of the eyes, which tends to occur when we look toward the horizon and can cause an increase in perceived size; Plug & Ross, 1994). Just as many different sources of depth information work together to create our impression of depth, many different factors may work together to create the moon illusion, and perhaps the other illusions as well.

SOMETHING TO CONSIDER:

Depth Information Across Species

Humans make use of a number of different sources of depth information in the environment. But what about other species? Many animals have excellent depth perception. Cats leap on their prey; monkeys swing from one branch to the next; and a male housefly maintains a constant distance of about 10 cm as it follows a flying female. There is no doubt that many animals are able to judge distances in their environment, but what depth information do they use? Considering the information used by different animals, we find that animals use the entire range of cues described in this chapter. Some animals use many cues, and others rely on just one or two.

To make use of binocular disparity, an animal must have eyes that have overlapping visual fields. Thus, animals such as cats, monkeys, and humans that have **frontal eyes** (Figure 10.46a), which result in overlapping fields of view, can use disparity to perceive depth. Animals with **lateral eyes**, such as the rabbit (Figure 10.46b), do not have overlapping visual fields and therefore cannot use disparity to perceive depth. Note, however, that in sacrificing binocular disparity, animals with lateral eyes gain a wider field of view—something that is extremely important for animals that need to constantly be on the lookout for predators.

The pigeon is an example of an animal with lateral eyes that are placed so that the visual fields of the left and right eyes overlap only in a 35-degree area surrounding the pigeon's beak. This overlapping area, however, happens to be exactly where pieces of grain would be located when the pigeon is



Bruce Goldstein

(a)



Barbara Goldstein

(b)

Figure 10.46 (a) Frontal eyes, such as those of the cat, have overlapping fields of view that provide good depth perception. (b) Lateral eyes, such as those of the rabbit, provide a panoramic view but poorer depth perception.

pecking at them, and psychophysical experiments have shown that the pigeon does have a small area of binocular depth perception right in front of its beak (McFadden, 1987; McFadden & Wild, 1986).

Movement parallax is probably insects' most important method of judging distance, and they use it in a number of different ways (Collett, 1978; Srinivasan & Venkatesh, 1997). For example, the locust uses a "peering" response—moving its body from side to side to create movement of its head—as it observes potential prey. T. S. Collett (1978) measured a locust's "peering amplitude"—the distance of this side-to-side sway—as it observed prey at different distances, and found that the locust swayed more when targets were farther away. Since more distant objects move less across the retina than nearer objects for a given amount of observer movement (Figure 10.10), a larger sway would be needed to cause the image of a far object to move the same distance across the retina as the image of a near object. The locust may therefore be judging distance by noting how much sway is needed to cause the image to move a certain distance across its retina (also see Sobel, 1990).

These examples show how depth can be determined from different sources of information in light. But bats, some of which are blind to light, use a form of energy we usually associate with sound to sense depth. Bats sense objects by using a method similar to the sonar system used in World War II to detect underwater objects such as submarines and mines. Sonar, which stands for **sound navigation and ranging**, works by sending out pulses of sound and using information contained in the echoes of this sound to determine the location of objects. Donald Griffin (1944) coined the term **echolocation** to describe the biological sonar system used by bats to avoid objects in the dark.

Bats emit pulsed sounds that are far above the upper limit of human hearing, and they sense objects' distances by noting the interval between when they send out the pulse and when they receive the echo (Figure 10.47). Since they use sound echoes to sense objects, they can avoid obstacles even when it is totally dark (Suga, 1990). Although we don't have any way of knowing what the bat experiences when these



Figure 10.47 When a bat sends out its pulses, it receives echoes from a number of objects in the environment. This figure shows the echoes received by the bat from (a) a nearby moth; (b) a tree located about 2 meters away; and (c) a house located about 4 meters away. The echoes from more distant objects take longer to return. The bat locates the positions of objects in the environment by sensing how long it takes the echoes to return. © Cengage Learning

echoes return, we do know that the timing of these echoes provides the information the bat needs to locate objects in its environment. (Also see von der Emde et al., 1998, for a description of how electric fish sense depth based on "electrollocation.") From these examples, we can see that animals use a number of different types of information to determine depth, with the type of information used depending on the animal's specific needs and on its anatomy and physiological makeup.

DEVELOPMENTAL DIMENSION: Infant Depth Perception

At what age are infants able to use different kinds of depth information? The answer to this question is that different types of information become operative at different times. Binocular disparity becomes functional early, and pictorial depth cues become functional later.

Using Binocular Disparity

One requirement for the operation of binocular disparity is that the eyes must be able to **binocularly fixate**, so that the two eyes are both looking directly at the object and the two

foveas are directed to exactly the same place. Newborns have only a rudimentary, imprecise ability to fixate binocularly, especially on objects that are changing in depth (Slater & Findlay, 1975).

Richard Aslin (1977) determined when binocular fixation develops by making some simple observations. He filmed infants' eyes while he moved a target back and forth between 12 cm and 57 cm from the infant. When the infant is directing both eyes at a target, the eyes should diverge (rotate outward) as the target moves away and should converge (rotate inward) as the target moves closer. Aslin's films indicate that although some divergence and convergence do occur in 1- and 2-month-old

infants, these eye movements do not reliably direct both eyes toward the target until about 3 months of age.

Although binocular fixation may be present by 3 months of age, this does not guarantee that the infant can use the resulting disparity information to perceive depth. To determine when infants can use this information to perceive depth, Robert Fox and coworkers (1980) presented random-dot stereograms to infants ranging in age from 2 to 6 months (see page 242 to review random-dot stereograms).

The beauty of random-dot stereograms is that the binocular disparity information in the stereograms results in stereopsis. This occurs only (1) if the stereogram is observed with a device that presents one picture to the left eye and the other picture to the right eye and (2) if the observer's visual system can convert this disparity information into the perception of depth. Thus, if we present a random-dot stereogram to an infant whose visual system cannot yet use disparity information, all he or she sees is a random collection of dots.

In Fox's experiment, an infant wearing special viewing glasses was seated in his or her mother's lap in front of a television screen (Figure 10.48). The child viewed a random-dot stereogram that appeared, to an observer sensitive to disparity information, as a rectangle-in-depth, moving either to the left or to the right. Fox's premise was that an infant sensitive to disparity will move his or her eyes to follow the moving rectangle. He found that infants younger than about 3 months of age would not follow the rectangle, but that infants between 3 and 6 months of age would follow it. He therefore concluded that the ability to use disparity information to perceive depth emerges sometime between 3½ and 6 months of age. This time for the emergence of binocular depth perception has been confirmed by other research using a variety of different methods (Held, Birch, & Gwiazda, 1980; Shimojo et al., 1986; Teller, 1997).

Another type of depth information is provided by pictorial cues. These cues develop later than disparity, presumably

because they depend on experience with the environment and the development of cognitive capabilities. In general, infants begin to use pictorial cues such as overlap, familiar size, relative size, shading, linear perspective, and texture gradients sometime between about 5 and 7 months of age (Kavšek, Granrud, & Yonas (2009); Yonas et al., 1982). We will describe research on two of these cues: familiar size and cast shadows.

Depth From Familiar Size

Granrud, Haake, and Yonas (1985) conducted a two-part experiment to see whether infants can use their knowledge of the sizes of objects to help them perceive depth. In the *familiarization period*, 5- and 7-month-old infants played with a pair of wooden objects for 10 minutes. One of these objects was large (Figure 10.49a), and one was small (Figure 10.49b). In the *test period*, which occurred about a minute after the familiarization period, objects (c) and (d) were presented at the same distance from the infant. The prediction was that infants sensitive to familiar size would perceive the object at (c) to be closer if they remembered, from the familiarization period, that this shape was smaller than the other one. In other words, if the infant remembered the green object as being small, then seeing it as big in their field of view could lead the infant to think it was the same small object, but located much closer. How can we determine whether an infant perceives one object as closer than another? The most widely used method is observing an infant's reaching behavior.

METHOD

Preferential Reaching

The preferential reaching procedure is based on observations that infants as young as 2 months old will reach for nearby objects and that 5-month-old infants are extremely likely to reach for an object that is placed within their reach and unlikely to reach for an




Figure 10.48 The setup used by Fox et al. (1980) to test infants' ability to use binocular disparity information. If the infant can use disparity information to see depth, he or she sees a rectangle moving back and forth in front of the screen. Adapted from "Assessment of Stereopsis in Human Infants," by S. L. Shea, R. Fox, R. Aslin, & S. T. Dumais, 1980, *Investigative Ophthalmology and Visual Science*, 19, 1440–1404, figure 1. Copyright © 1980, with permission from Elsevier.



object that is beyond their reach (Yonas & Hartman, 1993). Infant's sensitivity to depth has therefore been measured by presenting two objects side by side. As with the preferential looking procedure (Chapter 2, page 46), the left-right position of the objects is changed across trials. The ability to perceive depth is inferred when the infant consistently reaches more for the object that contains information indicating it is closer. When a real depth difference is presented, infants use binocular information and reach for the closer object almost 100 percent of the time. To test infants' use of pictorial depth information only, an eye patch is placed on one eye (this eliminates the availability of binocular information, which overrides pictorial depth cues). If infants are sensitive to the pictorial depth information, they reach for the apparently closer object approximately 60 percent of the time.

When Granrud and coworkers presented the objects to infants, 7-month-old infants reached for object (c), as would be predicted if they perceived it as being closer than object (d). The 5-month-olds, however, did not reach for object (c), which indicated that these infants did not use familiar size as information for depth. Thus, the ability to use familiar size to perceive depth appears to develop sometime between 5 and 7 months.

This experiment is interesting not only because it indicates when the ability to use familiar size develops, but also because the infant's response in the test phase depends on a cognitive ability—the ability to remember the sizes of the objects that he or she played with in the familiarization phase. The 7-month-old infant's depth response in this situation is therefore based on both what is perceived and what is remembered.

Depth From Cast Shadows

We know that shadows provide information indicating an object's position relative to a surface, as occurred in Figure 10.8. To determine when this ability is present in infants, Albert Yonas and Carl Granrud (2006) presented 5- and 7-month-old infants with a display like the one in **Figure 10.50**. Adults and older children consistently report

that the object on the right appears nearer than the object on the left. When the infants viewed this display monocularly (to eliminate binocular depth information that would indicate that the objects were actually flat), the 5-month-old infants reached for both the right and left objects on 50 percent of the trials, indicating no preference for the right object. However, the 7-month-old infants reached for the right object on 59 percent of the trials. Yonas and Granrud concluded from this result that 7-month-old infants perceive depth information provided by cast shadows.

This finding fits with other research that indicates that sensitivity to pictorial depth cues develops between 5 and 7 months. But what makes these results especially interesting is that they imply that the infants were able to tell that the dark areas under the toy were shadows and not dark markings on the wall. It is likely that this ability, like the other pictorial depth cues, is based largely on learning from interacting with objects in the environment. In this case, infants need to know something about shadows, including an understanding that most light comes from above (see page 111). **VL**

TEST YOURSELF 10.2

1. Describe the Holway and Boring experiment. What do the results of this experiment tell us about how size perception is influenced by depth perception?
2. What are some examples of situations in which our perception of an object's size is determined by the object's visual angle? Under what conditions does this occur?
3. What is size constancy, and under what conditions does it occur?
4. What is size-distance scaling? How does it explain size constancy?
5. Describe two types of information (other than depth) that can influence our perception of size.
6. Describe how illusions of size, such as the Müller-Lyer illusion, the Ponzo illusion, the Ames room, and the moon illusion, can be explained in terms of size-distance scaling.




Figure 10.50 Stimuli presented to 5- and 7-month-old children in Yonas and Granrud's (2006) cast shadow experiment. From Yonas, A., & Granrud, C. E. (2006). Infants' perception of depth from cast shadows. *Perception and Psychophysics*, 68, 154–160. Reproduced by permission.

7. What are some problems with the size–distance scaling explanation of (a) the Müller-Lyer illusion and (b) the moon illusion? What alternative explanations have been proposed?
8. Describe experiments that showed when infants can perceive depth using binocular disparity and using pictorial (monocular) cues. Which develops first? What methods were used?

THINK ABOUT IT

1. One of the triumphs of art is creating the impression of depth on a two-dimensional canvas. Go to a museum or look at pictures in an art book, and identify the depth information that helps increase the perception of depth in these pictures. You may also notice that you perceive less depth in some pictures, especially abstract ones. In fact, some artists purposely create pictures that are perceived as “flat.” What steps do these artists have to take to accomplish this? (p. 231)
2. Texture gradients are said to provide information for depth perception because elements in a scene become more densely packed as distance increases. The examples of texture gradients in Figures 10.4 and 10.7 contain regularly spaced elements that extend over large distances. But regularly spaced elements are more the exception than the rule in the environment. Make an informal survey of your environment, both inside and outside, and decide (a) whether texture gradients are present in your environment and (b) if you think the principle behind texture gradients could contribute to the perception of depth even if the texture information in the environment is not as obvious as the examples in this chapter. (p. 233)
3. How could you determine the contribution of binocular vision to depth perception? One way would be to close one eye and notice how this affects your perception. Try this, and describe any changes you notice. Then devise a way to quantitatively measure the accuracy of depth perception that is possible with two-eyed and one-eyed vision. (p. 236)

KEY TERMS

- | | | |
|---|--|--|
| Absolute disparity (p. 237) | Disparity tuning curve (p. 241) | Perspective convergence (p. 230) |
| Accretion (p. 232) | Echolocation (p. 254) | Pictorial cue (p. 229) |
| Active method (3-D TV) (p. 239) | Emmert’s law (p. 247) | Ponzo illusion (p. 250) |
| Ames room (p. 251) | Familiar size (p. 230) | Random-dot stereogram (p. 240) |
| Angle of disparity (p. 237) | Frontal eyes (p. 253) | Relative disparity (p. 238) |
| Angular size contrast theory (p. 253) | Horopter (p. 237) | Relative height (p. 229) |
| Apparent distance theory (p. 252) | Lateral eyes (p. 253) | Relative size (p. 229) |
| Atmospheric perspective (p. 230) | Lenticular projection (p. 240) | Size constancy (p. 246) |
| Binocular depth cell (p. 241) | Misapplied size constancy scaling (p. 248) | Size–distance scaling (p. 247) |
| Binocular disparity (p. 236) | Monocular cue (p. 228) | Stereopsis (p. 238) |
| Binocularly fixate (p. 254) | Moon illusion (p. 252) | Stereoscope (p. 238) |
| Conflicting cues theory (p. 250) | Motion parallax (p. 232) | Stereoscopic depth perception (p. 234) |
| Correspondence problem (p. 241) | Müller-Lyer illusion (p. 248) | Stereoscopic vision (p. 234) |
| Corresponding retinal points (p. 236) | Noncorresponding points (p. 237) | Strabismus (p. 235) |
| Cue approach to depth perception (p. 228) | Occlusion (p. 228) | Texture gradient (p. 231) |
| Deletion (p. 232) | Oculomotor cue (p. 228) | Visual angle (p. 244) |
| Disparity-selective cell (p. 241) | Passive method (3-D TV) (p. 239) | |

MEDIA RESOURCES

CourseMate

Go to CengageBrain.com to access Psychology CourseMate, where you will find the Virtual Labs plus an interactive eBook, flashcards, quizzes, videos, and more.

Virtual Labs

The Virtual Labs are designed to help you get the most out of this course. The Virtual Lab icons direct you to specific media demonstrations and experiments designed to help you visualize what you are reading about. The numbers below indicate the number of the Virtual Lab you can access through Psychology CourseMate.

10.1 Convergence (p. 228)

Description of how the eyes converge when looking at near objects.

10.2 Shape From Shading (p. 232)

Description of how shading facilitates perception for a three-dimensional object.

10.3 Ball in a Box (p. 232)

Computer animation showing how shadows affect perception of the location of a ball rolling in a box. (Courtesy of Daniel Kersten)

10.4 Illusory Motion in Depth (p. 232)

How a moving shadow can make a square appear to move in depth. (Courtesy of Daniel Kersten)

10.5 Size Perception and Depth (p. 251)

How perspective cues influence size perception.

10.6 Size Constancy in Infants (p. 256)

Albert Yonas's research on depth perception in infants.

- Tritanopia, 209, 210
- Tuning curves
 - disparity, 241
 - frequency, 277, 279
 - orientation, 64, 65
- Tunnel vision, 28
- 2-deoxyglucose (2DG) technique, 377
- Two-flash illusion, 311
- Two-point threshold, 343–344
- Tympanic membrane, 271, 272
- Umami tastes, 364, 368
- Unattended stimuli, 140
- Unconscious inference, 113
- Uniform connectedness, principle of, 104
- Unilateral dichromats, 209
- Univariance, principle of, 206
- Vagus nerve, 365
- Valid trials, 133
- Ventral pathway, 84, 164
- Ventral posterior nucleus, 349–350
- Ventriloquism effect, 311
- Ventrolateral nucleus, 340
- Vestibular system, 309
- Vibration perception, 345–346
- Video microscopy, 370
- Viewpoint invariance, 99–100
- Virtual Lab icons, 3, 19. *See also* Media resources
- Virtual museum experiment, 161
- Visible light, 22, 23, 201
- Visible spectrum, 201
- Vision
 - attention and, 127–150
 - balance and, 156–157
 - binocular rivalry and, 115–116
 - color perception and, 199–225
 - computer-based systems and, 97
 - dark adaptation and, 27, 29–32
 - depth perception and, 227–243
 - detail and, 43–44
 - development of, 46–47
 - eye movements and, 128–129
 - feature detectors and, 66–70, 81–83
 - focusing process and, 23–24
 - hearing related to, 310–313
 - impairments of, 24–26, 28–29, 32
 - infant perception and, 45–47
 - location information for, 290, 291
 - memory and, 89–91
 - motion perception and, 175–197
 - myopia and, 25–26
 - navigation and, 157–163
 - object perception and, 95–124
 - olfaction compared to, 374
 - persistence of, 109
 - reaching/grasping and, 163–166
 - scene perception and, 81–83, 108–114, 117–119
 - size perception and, 243–253
 - spectral sensitivity and, 32–34
 - speech perception and, 324
 - steps in process of, 22
 - stereoscopic, 234
 - stimulus for, 22
 - taste compared to, 369
 - See also* Visual system
- Visual acuity, 43–47
 - development of, 46–47
 - rods vs. cones and, 43–44, 47
 - test of, 45–46
- Visual angle, 244–246
- Visual attention. *See* Attention
- Visual capture, 311
- Visual direction strategy, 159
- Visual evoked potential (VEP), 46
- Visual form agnosia, 8
- Visual illusions, 248–253
 - Ames room, 251
 - moon illusion, 252–253
 - Müller-Lyer illusion, 248–250
 - Ponzo illusion, 250–251
 - size illusion, 86
 - waterfall illusion, 179
 - White's illusion, 59–60
- Visual impairments
 - age-related, 24–25
 - color blindness, 199–200, 209
 - color deficiency, 208–210
 - detached retina, 32
 - focusing problems, 25–26
 - macular degeneration, 28
 - retinitis pigmentosa, 28, 29
- Visualizing
 - colors, 211
 - scenes/objects, 112
- Visual masking stimulus, 109
- Visual perception
 - demonstrations of, 10, 24, 29
 - feature detectors and, 66–69
 - infants and, 45–47
 - neural processing and, 40–44, 53–72
 - perceptual organization and, 77–78
- Visual pigments, 7, 23
 - bleaching of, 31–32
 - color perception and, 205, 206–207, 210
 - limiting function of, 44–45
 - molecules of, 26–27
 - regeneration of, 31–32
 - spectral sensitivity and, 33
- Visual receiving area, 35, 64
- Visual receptors, 6–7, 26–34
- Visual scanning, 128
- Visual search, 17, 144–145
- Visual system
 - color perception and, 204–207, 212–213
 - columnar organization of, 80–81
 - convergence in, 41–44
 - development of, 46–47
 - feature detectors in, 66–70
 - illustrated overview of, 63
 - impairments of, 24–26, 28–29, 32
 - lateral geniculate nucleus and, 63–64
 - lateral inhibition in, 54–59
 - magnification factor in, 78–80
 - modularity and, 87–89
 - myopia and, 25–26
 - neural processing in, 40–44, 53–72
 - receptors in, 6–7, 26–34
 - retinotopic maps and, 78
 - review questions on, 34–35
 - scene perception and, 81–83
 - spatial organization in, 78–83
 - streams for information in, 83–86
 - striate cortex and, 64–66
 - transduction process in, 26–27
 - See also* Eyes; Vision
- Visual transduction, 26–27
- Visuomotor grip cells, 164
- Vocal characteristics, 328
- Vocal tract, 318
- Voice cells, 329
- Voice onset time (VOT), 322–324, 332
- Voice recognition, 284–285
- Vowels, 318, 319, 320
- Voxels, 117
- Walking experiments, 159–160
- Waterfall illusion, 179
- Wavelength, 22, 23, 201–204, 221
- Waves* (Hurskainen), 103
- Wayfinding, 160–163
 - effect of brain damage on, 162–163
 - importance of landmarks in, 160–162
- Weber fraction, 15
- Weber's law, 15
- Web resources. *See* Media resources
- Wernicke's aphasia, 329
- Wernicke's area, 329
- What* pathway
 - for audition, 299–300
 - for vision, 84, 86, 164
- Where* pathway
 - for audition, 299–300
 - for vision, 84, 164
- Whiteout conditions, 243
- White's illusion, 59–60
- Who's Afraid of Virginia Woolf?* (film), 145, 146
- Winter Twilight* (Baker), 122
- Word deafness, 329
- Word perception, 325–328
 - breaks between words, 326–328
 - infant development and, 327–328
 - speaker characteristics and, 328
 - words in sentences, 325–326
- Young-Helmholtz theory of color vision, 204
- Zombies in the brain, 393

