

01 - 5 Perception

5 Perception

CHAPTER 5 PERCEPTION © ISTOCKPHOTO.COM/JUSTIN HORROCKS For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

On a warm Saturday some years ago, two young men – we’ll call them Alex and Simon – left their homes for a day’s hunting trip. As they walked along an abandoned road, their conversation touched on various hunting-related topics, but mostly they talked about bears. Alex had seen a bear the previous weekend, and both men were apprehensive about these dangerous creatures. They knew that their hunting rifles were powerful, but they were well aware that bears were equally powerful. The hunters maintained a constant vigil. It was almost midnight by the time Alex and Simon retraced their path along the road bound for home. There was no moon; the forest was quiet and dark. The two hunters were tired from their day’s efforts. As they rounded a curve, they suddenly became aware of a low growling sound which they perceived to come from a large, dimly illuminated animal quivering slowly but ominously in the middle of the road, about 50 meters away. Terrified, they raised their rifles and fired. The growling noise and the quivering abruptly ceased. An instant later an unmistakably human scream pierced the night. The hunters’ relief at having killed the bear was replaced by confusion and dismay as they realized that the bear wasn’t a bear at all. It was a tent in which had dwelt two campers. One of the campers now lay dead from a bullet wound, while the other knelt above him, wailing in horror. Investigation carried out in the aftermath of this terrible event revealed that Simon’s bullets had passed harmlessly through the tent; it was one of the bullets from Alex’s gun that killed the camper. Accordingly, Alex went to trial, accused of negligent homicide. The tragedy of the killing was mirrored in the courtroom by Alex’s overwhelming sorrow about what had happened. There was one critical fact, however, about which both Alex and Simon were certain: They had perceived a bear, not a tent that night. ‘We never would have shot if we had had any idea that it wasn’t a bear’, they both swore. The prosecutor dismissed these assertions as ridiculous and desperate lies: The bullet-riddled tent itself was placed in the center of the courtroom and the prosecutors asked the jury, ‘How could the defendant have possibly mistaken this rectangular yellow tent for a furry brown bear?’ prosecution’s question seems quite reasonable. There, sitting in the courtroom, for all to behold, was a big © The McNeil River State Game Sanctuary, Photo by Jon C. Pascal For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

CHAPTER OUTLINE

WHAT IS THE USE OF PERCEPTION? Processing and using incoming sensory information Five functions of perception

ATTENTION Selective attention Auditory attention Attention, perception, and memory Costs and benefits of selectively attending to stimuli

CUTTING EDGE RESEARCH: DISTRACTION VIA VIRTUAL REALITY DIMINISHES SEVERE PAIN LOCALIZATION Separation of objects Perceiving distance Perceiving motion

RECOGNITION Global-to-local processing The binding problem: pre-attentive and

attentive processes Determining what an object is Later stages of recognition: network models
Recognizing natural objects and top-down processing Special processing of socially relevant
stimuli: face recognition Failure of recognition ABSTRACTION Exact to abstract The advantages of
abstraction: required storage and processing speed PERCEPTUAL CONSTANCIES The nature of
constancies Color and brightness constancy Shape constancy Size constancy Illusions Constancies
in all sensory modalities DIVISIONS OF LABOR IN THE BRAIN The neural basis of attention The
visual cortex Recognition versus localization systems PERCEPTUAL DEVELOPMENT How indeed? On
the face of it, the Discrimination by infants Controlled stimulation SEEING BOTH SIDES: IS
PERCEPTUAL DEVELOPMENT AN INNATE OR SOCIALLY ACQUIRED PROCESS? 151

yellow tent, appearing not at all similar to a bear. However, a half-century's research on perception – visual perception in this instance – suggests that under the circumstances, it wasn't at all unreasonable for Simon and Alex to have perceived the tent to be a bear. In this chapter we will elaborate on why this is so, demonstrating in the process how the raw sensations that we discussed in Chapter 4 become translated into the perceptions that are directly responsible for our behavior. To get a feel for what we mean by this, let's start with a couple of demonstrations. Look first at the left panel of Figure 5.1. Do you recognize an object? If you are like most people (and have not seen this demonstration previously), your answer would be, 'No'. Now look at the right panel of Figure 5.1. What does it say? Again if you're normal and haven't seen this demonstration before, you probably read, 'I LOVE PARIS IN THE SPRINGTIME'. In both cases you had perceptions, of meaningless black-and-white blobs in one instance and of a common cliché in the other, that somehow derived from the basic, objective stimulus, that is the light that entered your eyes and fell on your retina. In both instances, however, there are interesting and systematic disconnects between the raw data and the ensuing perception. Does that 'I love Paris' statement really say what you thought it did? Look at it again, this time reading it very slowly and word-by-word. You will see that it actually says, 'I love Paris in the the Springtime'. What about those meaningless blobs in the left panel of Figure 5.1? Look at the picture on page 193 and return here when you've done so. Are you back? The left panel of Figure 5.1 is no longer meaningless, is it? Indeed, if you are like most people, it is difficult for you to believe that it ever was meaningless. The stimulus entering your eyes is identical to what it was before, but the perception is entirely different: The black-and-white blobs are now organized into a meaningful object. These demonstrations are designed to convince you that while information may enter our senses in bits and pieces, that is not how we perceive the world. We perceive a world of objects and people, a world that gracefully presents us with integrated wholes, rather than chaotically bombarding us with piecemeal sensations. Only under unusual circumstances, or when we are drawing or painting, do we notice the individual features and parts of stimuli; most of the time we see three-dimensional objects, hear words and music, taste and smell the frying fish and chips, and feel a hand on our arm. WHAT IS THE USE OF PERCEPTION? Any living organism must solve an unending series of problems presented to it by the environment within which it dwells. The complexity of the problems and associated sophistication of the solutions depend on the nature and complexity of the organism. If you are a daffodil, for example, the problems you must deal with are relatively simple. You must figure out where your roots should go on the basis of the soil structure you're planted in, determining in the process, the soil's texture, along with the distribution within the soil of moisture and nutrients. Additionally you must Figure 5.1 Raw Data and the Resulting Perception. Left panel: Do you see a meaningful object? (Look at Figure 5-38 on page 193 if you need help.) Right panel: What does the phrase say? CHAPTER 5 PERCEPTION For more Cengage Learning textbooks, visit

determine which way to orient yourself on the basis of where the sun is. But that's about it for daffodils. Humans, it won't surprise you to hear, are quite a bit more complex. With respect to perception, the most important differences between daffodils and humans are these: First a human is mobile: The vast majority of us must make our way through the environment, determining in the process, the potential routes that we could take and the obstacles that must be surmounted for each route. Second, a human manipulates objects: We turn the steering wheel on a car, sign our names with a pen, and kick a ball toward the goal. Third, a human makes decisions on the basis of symbols such as written or spoken words or hieroglyphics. Fourth, a human makes and executes complex plans to deal with sudden unexpected events: upon glimpsing a sinister form in a dark alley, we evaluate our options and cross to the other side of the street where we can seek safety in the crowd that has gathered there.

Processing and using incoming sensory information How do we do this? One possibility is that the information from the environment – in the case of vision, the environment's two-dimensional representation on our retina – is all that is really necessary to live a normal life. The American J. J. Gibson offered a theory of ecological optics, which specified just that. According to Gibson, the vast richness of optical information from the world – the change in texture with distance, the shifting of objects' images relative to one another as one walks by them, and so on – is sufficient to solve all vision-related problems that the world presents us. Although ingenious, sophisticated and useful, Gibson's theory has been rejected by most perception scientists as insufficient. Instead, it is argued, humans require a continually updated image or a model of the environment within our brains, and it is then based on that model that humans perceive, make decisions, and behave. Two ingredients are necessary to formulate and maintain such a model. The first is some means of acquiring raw information about the environment. In Chapter 4, we discussed how our sense organs are used to accomplish this. But acquiring raw information is not sufficient to build a model, any more than acquiring a stack of wood is sufficient to build a house. In addition, we need a means of organizing all this raw information into some kind of coherent structure. Such organization is not simple. Most basically, perception of the world involves solving what is referred to as the many-to-one problem. Illustrated in vision, this problem boils down to the mathematical necessity that many configurations of objects in the environment all give rise to the same representation on the retina. Later in this chapter we will have quite a bit more to say about this. For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

WHAT IS THE USE OF PERCEPTION? For the moment, to illustrate, think of seeing a pine tree in the distance. A 2-meter-high tree seen from a distance of 100 meters would produce the same-size retinal image as a 4-meter-high tree seen from a distance of 200 meters (as would an infinite number of other height-distance combinations). The many-to-one problem entails deciding, based on the one retinal image, which of the infinite possible size-distance configurations gives rise to the retinal image. The visual system must solve this problem by using other information – both information already stored in the brain (e.g., these trees are Christmas trees which are generally 2-meter rather than 4-meter trees) and additional visual cues (e.g., the person standing next to the tree is about the same height as the tree). More generally, making inferences from the sensory data back to the state of the environment that gave rise to the data requires assumptions about how the world is put together – birds are usually to be found above horses, stoves are usually to be found near refrigerators, a scene is usually illuminated by a single kind of light source, and so on. Thus, perception is the use of such assumptions to integrate incoming sensory information into a model of the world, based upon which we make decisions and take action. Usually this process works

pretty efficiently and, for example, a yellow tent in the environment produces a model – a perception – of a yellow tent in our mind. Sometimes it doesn't work so well: A yellow tent in the environment produces the perception of a bear in our mind, and we shoot it. Generally speaking, each sensory modality – seeing, hearing, and so on – has both a sense organ involved in acquiring the raw information from the environment and a more central system in the brain for transforming this information into organized percepts. Perception is sufficiently complex that any classification of it must be somewhat arbitrary. For organizational purposes, however, it is useful to divide perceptual issues into five categories. First via the process of attention a decision must be made about which incoming information is to be further processed, and which is to be discarded (should I be eavesdropping on the conversation on my left which seems to be about my spouse or the conversation on my right which seems to involve cricket scores?). Second, the system must be able to determine where objects of interest are (is that potentially dangerous object at arm's length, on my left, hundreds of meters straight ahead, or where?). Third, the perceptual system must be able to determine which objects are out there (is that a tent or a bear that I'm looking at?). Fourth, the system must be able to abstract the critical features of a recognized object (a couch that has wrinkles and bumps in it would be reasonably perceived and described as 'rectangular' even though its shape isn't a perfect rectangle). This abstraction ability is closely related to the

154 CHAPTER 5 PERCEPTION fifth category of perceptual issues, that of perceptual constancy: The perceptual system must maintain certain inherent features of objects (e.g., a door's inherent rectangular shape) even when the door's angle to you is such that it forms a trapezoid on your retina. In the next five sections, we will discuss these five issues: attention, localization, recognition, abstraction, and constancy. We will then discuss some of the biological correlates of these perceptual processes. Finally, we consider the development of perception. Throughout the chapter we focus primarily on visual perception because this is the area that has been most investigated. Keep in mind, though, that the goals of localization, recognition, and constancy apply to all sensory modalities. With regard to recognition, for example, we can use our hearing to recognize a Mozart sonata, our sense of smell to recognize fish and chips, our sense of touch to recognize our keys in our trouser pocket, and our body senses to recognize that we are upright in a dark room. INTERIM SUMMARY I The study of perception deals with the question of how organisms process and organize incoming raw sensory information in order to (1) form a coherent representation or model of the world within which the organism dwells and (2) use that representation to solve naturally occurring problems, such as navigating, grasping, and planning. I Five major functions of the perceptual system are: (1) determining which part of the sensory environment to attend to, (2) localizing, or determining where objects are, (3) recognizing, or determining what objects are, (4) abstracting the critical information from objects, and (5) keeping the appearance of objects constant, even though their retinal images are changing. Another area of study is how our perceptual capacities develop. ATTENTION We began the previous chapter, Sensory Processes, by underscoring that at any given instant our sense organs are being bombarded with a vast amount of information from the environment. As you sit reading, stop for a moment and attend to the various stimuli that are reaching you. There is, in your visual field, more than just the pages of this book. Perhaps your left shoe is feeling a little tight. What sounds do you hear? What odors are there in the air? Meanwhile the human bombardee is generally engaged in trying to accomplish some task. This task could be as simple as drinking a cup of coffee or as complex as doing For more Cengage Learning textbooks, visit www.cengagebrain.co.uk brain

surgery, or something in between like trying to digest the information in this book. Whatever the task, however, only a tiny portion of the incoming stream of information is relevant to it; the vast majority is irrelevant. This state of affairs implies that the sensory systems and the brain must have some means of screening the incoming information – allowing people to select only the information relevant to the task at hand for perceptual processing, and to ignore the irrelevant information. If such a screening process did not exist, the irrelevant information would overwhelm the relevant information, and we would never get anything done. The ability to selectively attend only to a small subset of all of the information in the environment is the topic of this section. This seemingly simple ability is now widely believed to involve three separate sets of processes that are anatomically distinct in the brain (e.g., Fan et al., 2002). One is responsible for keeping us alert. For example, an air-traffic controller needs to remain alert in order to remain aware of the various aircraft that she is responsible for; failure of this system might lead to a disastrous attentional lapse. A second system is responsible for orienting processing resources to task-relevant information (e.g., focusing on the voice so that we can understand what is being said), and the third, sometimes referred to as the ‘executive’, decides whether we want to continue attending to the information or instead switch attention to other information (e.g., ‘This person is talking about chloroplasts – I have no interest in chloroplasts’). The point is that rather than being a single process, attention is best thought of as involving multiple interacting processes. We describe these processes in more detail below.

Selective attention How exactly do we direct our attention to objects of interest? The simplest means is by physically reorienting our sensory receptors. For vision, this means moving our eyes until the object of interest falls on the fovea which, you will recall from Chapter 4, is the most sensitive region of the retina – the region designed to process visual detail.

Eye movements Studies of visual attention often involve observing an observer looking at a picture or scene. If we watch the person’s eyes, it is evident that they are not stationary. Instead visual scanning takes the form of brief periods during which the eyes are relatively stationary, called eye fixations, separated by quick jumps of the eye called saccades. Each fixation lasts approximately 300 milliseconds (about a third of a second) while saccades are very fast (on the order of 20 milliseconds). It is during the fixation periods that visual information is acquired from the environment; vision is essentially suppressed during saccades.

Figure 5.2 Eye Movements in Viewing a Picture. Next to the picture of the girl is a record of the eye movements made by an individual inspecting the picture. (D. L. Yarbuss (1967) ‘Eye Movements and Vision’, Plenum Publishing Corporation. Reproduced by permission of the publisher) By monitoring a person’s eye fixation pattern over a scene, we can gain considerable insight about the sequencing of the person’s visual attention. There are a number of techniques for recording eye movements but all of them eventually produce a millisecond-to-millisecond computer record of where on the scene is the gaze. Such a record can be used, among other things, to reproduce the scene itself along with the sequence of fixations on it, as shown in Figure 5.2. Generally speaking the points on which the eyes fixate are not random, but rather are the areas of the scene that contain the most information. The exact definition of ‘information’ is beyond the scope of this book, but in this context it refers roughly to those areas that are most likely to distinguish the scene being viewed from any other similar scene. For example, as shown in Figure 5.2, a person looking at a face makes many fixations on the eyes, nose, and mouth – those features that most efficiently distinguish one face from another. Loftus and Mackworth (1978) demonstrated the relation between fixations and pictorial information by presenting a picture containing an object that was either unusual or not unusual within some background context. For instance, one observer might

be shown a picture of a farmyard with a tractor in the middle of it, while another observer would see the same farmyard picture but with an octopus rather than the tractor. Eye fixations were directed earlier and more frequently to the unusual object (the octopus) than to the normal object (the tractor). (For control purposes, two other observers would see pictures with, respectively, an octopus in an underwater scene and a tractor in the same underwater scene; here the tractor would be the unusual object and the octopus would be the normal object). For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

ATTENTION

Weapon focus A useful practical application of this kind of eye movement research concerns what is referred to as weapon focus: Victims of armed crimes are often able to very accurately describe what the weapon looked like, but seem to know relatively little about other aspects of the scene, such as the appearance of the person who was wielding the weapon, suggesting that attention was primarily focused on the weapon. Laboratory studies have generally confirmed this anecdotal evidence (see Steblay, 1992). Loftus, Loftus, and Messo (1987) recorded eye movements while observers looked at a slide sequence, one of which showed a person handling a critical object which was either benign (a checkbook) or threatening (a knife). They found that more eye fixations occurred on the critical object compared to the rest of the scene when the object was threatening than when it was benign; correspondingly, observers were less able to recognize other aspects of the scene, such as the face of the person holding the object, when they had viewed a threatening compared to a benign object. It's important to note that the laboratory studies undoubtedly underestimate the attention-demanding power of a weapon compared to the real-life situations that they are meant to explore. In both the real-life and the laboratory situations, a weapon is unusual and would be expected to draw attention on that basis, as described above. However, the real-life situation has the added component that the weapon constitutes crucial environmental information relevant to what becomes the threatened individual's immediate task: that of survival.

Directed attention without eye movements Although we normally attend to what our eyes are pointed at, we can also selectively attend to a visual stimulus without moving our eyes. In experiments that demonstrate this, observers have to detect when an object occurs. On each trial, the person stares at a blank field, then sees a brief cue directing them to attend either to the left or to the right. An object is then presented either in the location indicated by the cue or in the opposite location. The interval between the cue and object is too brief for observers to move their eyes, yet they can detect the object faster when it occurs in the cued location than elsewhere. Presumably, they are attending to the cued location even though they cannot move their eyes there (Posner & Raichle, 1994).

Auditory attention Attention is multimodal; that is, attention can move within a modality (e.g., from one visual stimulus to another) or between modalities (we have all had the experience of shifting our attention from watching the road while driving to listening to the person who just called our cell phone). Much of the original research on

156 CHAPTER 5 PERCEPTION ^a ISTOCKPHOTO.COM/RENEE LEE Although we may hear a number of conversations around us, as at a cocktail party, we remember very little of what we do not attend to. This is known as selective listening. attention was done on auditory attention (e.g., Cherry, 1953). A real-life analogue of Cherry's work is a crowded party. The sounds of many voices bombard our ears. However, we can use purely mental means to selectively attend to the desired message. Some of the cues that we use to do this are the direction the sound is coming from, the speaker's lip movements, and the particular characteristics of the speaker's voice (pitch and intonation). Even in the absence of any of these cues, we can (though with difficulty) select one of two messages to follow on the basis of its meaning. Attention, perception, and memory With some

caveats to be described in Chapter 8, a general rule has emerged about the relation between attention and later memory: We are consciously unaware of, and remember little, if anything, about nonattended information. In the auditory domain, a procedure known as shadowing is used to demonstrate this. The observer wears stereo earphones; however, entirely different messages are played to the two different ears. The person is asked to repeat (or 'shadow') one of the messages as it is heard. After a few minutes the messages are turned off and the listener is asked about the unshadowed message. The listener's report of the message is usually limited to the physical characteristics of the sound in the unshadowed ear – whether the voice was high or low, male or female, and so forth; he or she can say almost nothing about the content of the message and, indeed, does not even notice when the language changes from English to French and then back again (Moray, 1969). Loftus (1972) reports an analogous finding in vision. He showed two pictures, side-by-side, but asked the observer to look at only one of them (and monitored the observer's eye movements to ensure compliance). The finding was that later memory was considerable for the attended picture, but was nil for the unattended picture. For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

The fact that we can report so little about auditory messages that we do not attend to initially led researchers to the idea that nonattended stimuli are filtered out completely (Broadbent, 1958). However, there is now considerable evidence that our perceptual system processes nonattended stimuli to some extent (in vision as well as audition), even though those stimuli rarely reach consciousness. One piece of evidence for partial processing of nonattended stimuli is that we are very likely to hear the sound of our own name, even when it is spoken softly in a nonattended conversation. This could not happen if the entire nonattended message were lost at lower levels of the perceptual system. Hence, lack of attention does not block messages entirely; rather, it attenuates them, much like a volume control that is turned down but not off (Treisman, 1969).

Costs and benefits of selectively attending to stimuli

As the previous section indicates, one cost of selectively attending to information is that observers are often oblivious to other, potentially important, stimuli in the environment. For example, Simons and Chabris (1999) showed participants a film of several students passing a basketball to one another; the observers' task was to count the total number of passes. During the film a person dressed in a gorilla suit slowly walked right through the middle of the scene. Because participants attended to the basketball almost nobody noticed the gorilla! This inattention blindness is closely related to change blindness, which is the failure of people to notice even large-scale changes to scenes. An interesting case of this was demonstrated by Simons and Levin (1998) on the campus of Cornell University New York State. In each trial of their experiment a student stopped a pedestrian to ask directions to a building. While the pedestrian responded, two people carrying an opaque door walked between the two people, temporarily blocking the pedestrian's view of the student; during this time the student switched places with one of the door carriers. Subjects noticed less than half of the time that they were now talking to a completely different person! Manipulations that drew attention to the speaker's face substantially reduced this change-blindness effect. That people can switch attention between sets of information has been recently put to interesting use by medical science in surgery for cataracts, which occur when the lens of the eye becomes cloudy so that it no longer adequately transmits light. The typical procedure is to remove the cloudy lens, replacing it with a clear artificial one. However, unlike a natural lens which can adjust its thickness to focus on objects at varying distances, artificial lenses are usually rigid. As a result, people who receive them can clearly see objects that are at least three feet away, but need special glasses to focus on

close objects and to read. New artificial lenses have been developed that consist in a set of numerous concentric rings, where alternating rings focus on close and far objects. As a result, two images are simultaneously projected onto the retina - one in which near objects are in focus and far ones are blurry, and a second where far but not near objects are in focus. Research indicates that patients who receive these lenses can selectively attend to one image or the other, and are unaware of the nonattended image. Thus a single fixed lens can provide clear perception for objects both near and far (e.g., Brydon, 2003).

INTERIM SUMMARY | Selective attention is the process by which we select some stimuli for further processing while ignoring others. In vision, the primary means of directing our attention are eye movements. Most eye fixations are on the more informative, i.e., unusual, parts of a scene. | Selective attention also occurs in audition. Usually we are able to selectively listen by using cues such as the direction from which the sound is coming and the voice characteristics of the speaker. | For the most part, we can only remember what we attend to. Our ability to selectively attend is mediated by processes that occur in the early stages of recognition as well as by processes that occur only after the message's meaning has been determined. | By not attending to - i.e., ignoring - large parts of the environment, we lose the ability to remember much about those parts of the environment. However, such selective attention pares down the amount of necessary information processing to the point where it is manageable by the brain.

CRITICAL THINKING QUESTIONS 1 It seems quite clear that attention can be monitored by watching where a person looks. Suppose that you hypothesize that selective visual attention could go from one place to another in the environment even with the eyes held still. How would you test this hypothesis? 2 How does selective attention aid perception under everyday circumstances? What would be the consequences of driving a car in a city where no one had the ability to attend selectively? What kinds of accidents might occur more frequently than occur now? Would any kinds of accidents be apt to occur less frequently? For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

LOCALIZATION LOCALIZATION Earlier, we described various problems that humans must solve for which localization of information is relevant. The most important such problems are (1) navigating our way around the often cluttered environment (think about what is required just to make your way from your bed to your kitchen sink without running into anything) and (2) grasping an object (to smoothly guide your fingers in quest of picking up your pen, you must know accurately where the pen is to begin with). To know where the objects in our environment are, the first thing that we have to do is separate the objects from one another and from the background. Then the perceptual system can determine the position of the objects in a three-dimensional world, including their distance from us and their patterns of movement. In this section we discuss each of these perceptual abilities in turn.

Separation of objects The image projected on our retina is a mosaic of varying brightnesses and colors. Somehow our perceptual system organizes that mosaic into a set of discrete objects projected against a background. This kind of organization was of great concern to Gestalt psychologists. (Recall from Chapter 1 that Gestalt psychology was an approach to psychology that began in Germany early in the twentieth century.) The Gestalt psychologists emphasized the importance of perceiving whole objects or forms, and proposed a number of principles to explain how we organize objects.

Figure and ground The most elementary form of perceptual organization is that in a stimulus with two or more distinct regions, we usually see part of it as a figure and the rest as ground (or background). The regions seen as a figure contain the objects of interest - they appear more solid than the ground and appear in front of it. Figure 5.3a shows that figure-ground organization can be ambiguous. When you look at this pattern you might see a pair of silhouette faces gazing at each other, or you might see an ornate vase. The vase appears white against a black ground, whereas the faces are black

against a white ground. Notice that as you look at Figure 5.3b for a few moments, the two pattern organizations alternate in consciousness, demonstrating that the organization into figure and ground is in your mind, not in the stimulus. Notice, also, that the faces and the vase never appear together. You 'know' that both are possible, but you cannot 'see' both at the same time. Generally speaking, the smaller an area or a shape, the more likely it is to be seen as figure. This is demonstrated by comparing Figures 5.3a, b, and c. It is easier to see the

CUTTING EDGE RESEARCH Distraction via Virtual Reality Diminishes Severe Pain While driving down a road in Baghdad in a U.S. military patrol humvee convoy, 21-year old Mark Powers was badly burned when a terrorist's roadside bomb exploded up into his vehicle. Deep flash burns to his hands, arms, chest and thighs covered over 32 percent of his body, and required skin grafting. Although opioid pain killers helped reduce his pain as he lay motionless in his hospital bed, they were much less effective during wound-care procedures. While having his wounds cleaned, Mark, like most burn patients, continued to experience severe to excruciating pain as well as numerous unpleasant side effects from the drugs. In response, the patient was given virtual reality SnowWorld to help reduce excessive pain from his combat-related burn injury (Maani, Hoffman et al., 2008). In 1996, Dr. Hunter Hoffman, from the University of Washington's Human Interface Technology Laboratory and Dr. David Patterson from Seattle's Harborview Hospital Burn Center co-originated a new psychological pain control technique - one that relied on diverted attention in a virtualreality (VR) setting - to supplement the usual drugs. Diverting attention is particularly useful with burn pain. The reason for this is that pain perception has a strong psychological component. As described in Chapter 4, pain, like any sensory input, consists of a specific signal, in this case, a train of nerve impulses from pain receptors in the skin. However, as we discuss in this chapter, perception, which is the interpretation of sensory input, is not entirely determined by the sensory input. This potential disconnect between sensation and perception is particularly salient with pain: The same incoming pain signal can be interpreted as painful or not painful, depending on what the patient is thinking and doing. To explore what happens in someone's brain when they experience virtual reality analgesia, the researchers designed a unique magnet-friendly fiberoptic photonic VR goggle system so subjects could have the illusion of going inside SnowWorld while scientists measured their brain activity. Since fMRI brain scanners measure changes in brain activity, Hoffman, Richards, et al. (2004) attached a small medical hotplate to the foot of healthy volunteers, which delivered 30 seconds of pain p 30 seconds of no pain, six times. Participants reported feeling strong pain when the hotplate was hot, and their brains showed increased activity in five brain areas of the brain associated with pain perception. Interestingly, when these participants went into SnowWorld, they reported large reductions in pain even when the hotplate was on, and the amount of pain-related brain activity dropped 50 to 97 percent in all five brain 'regions of interest'. In other words, fMRI brain scans provided objective evidence that VR reduces pain, and early clues to how VR reduces pain (see Hoffman, 2004). These results can be interpreted within the context of what is known as a gate control theory of pain. The idea here is that higher order thought processes such as attentional distraction, can initiate feedback signals from the cortex to the spinal cord, thereby inhibiting the intensity of incoming pain signals. In other words, in addition to influencing the way patients interpret incoming pain signals, distraction may actually reduce the intensity of the incoming pain signals. The problem with burn patients is that, unable to rise from their beds during wound care, they are not generally able to interact with any sort of interesting, attention-attracting realworld environment. Enter VR, which allows the patient to enter any world imaginable without physically going anywhere. A VR computer set up in the hospital

room sends video output to two miniature LCD screens positioned in front of the patient's eyes using a specially designed helmet. Motion sensors track the patient's head position and feed this information into the computer. When the patient moves his or her head (e.g., looks up), the computer updates the artificial environment accordingly (e.g., changing the image from a virtual river to a virtual sky). These real-time changes in sensory input, in response to patients' actions, afford the illusion of actually being in the computer-generated environment. In principle, a person's

a) b) c) Figure 5.3 Reversible Figure and Ground. Three patterns in which either a white vase or a pair of black faces can be seen. Note that it is impossible to see both organizations at the same time, even though you know that both are possible percepts. When the white area is smaller (a), the vase is more likely to be seen; when the black area is smaller (c), the faces are more likely to be seen. CHAPTER 5 PERCEPTION For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

perception within VR can perfectly mimic the perception of a person within the real world (as spectacularly envisioned by the science-fiction writer, Neil Stephenson in his novel, *Snow Crash*). An incoming pain signal requires conscious attention to be perceived as pain. But being drawn into another world – one of virtual reality – drains a substantial amount of attentional resources, leaving less available to process pain signals. Thus, the attentional 'spotlight' that would normally be focused on the pain is lured instead into the virtual world. For many patients undergoing VR treatment, their pain – particularly the normally excruciating pain associated with the care and cleansing of their wounds – becomes little more than an annoyance, distracting them from their primary goal of exploring the virtual world. In a preliminary case study (Hoffman, Doctor, Patterson, Carrougher and Furness, 2000), two patients with severe burns went into a VR environment consisting of a virtual kitchen complete with countertops, a window looking out at a partly cloudy sky, cabinets and doors. Patients could perform actions – pick up a teapot, plate, toaster, plant, or frying pan – by inserting their cyberhand into the virtual object, and clicking a grasp button on their 3-D mouse. Each patient could pick up a virtual wiggly-legged spider or eat a virtual chocolate bar that possessed solidity, weight and taste, created via a mixedreality force feedback technique developed by Hoffman. The VR treatments showed a great deal of promise with these two initial patients. Patient 1 had five staples removed from a burn skin graft while playing Nintendo (a control condition), and six staples removed from the same skin graft while in VR. He reported dramatic reductions in pain in the VR compared to the Nintendo condition. Patient 2, even with more severe and extensive burns, showed the same pattern. Hoffman, Patterson and Carrougher (2000), have found additional support that VR reduces burn pain. Twelve severely burned patients reported substantial pain reduction during physical therapy when in VR compared to conventional treatment. In addition to distracting the patients, VR can likely be used to motivate patients to perform necessary but normally very painful stretching motions, using behavioral reinforcement techniques. For example, while playing in a VR game they could vase when the white area is smaller, and it is easier to see the faces when the black area is smaller (Weisstein & Wong, 1986). These figure-ground principles are not restricted to simple stimuli. As shown in Figure 5.4, they apply to quite complex pictures as well. It should be noted that, while vision is the most salient source of figure-ground relations, we can also perceive figure-ground relations in other senses. For example, we may hear the song of a bird against a background of outdoor noises, or the melody played by a violin against the harmonies of the rest of the orchestra.) Grouping of objects We see not only objects against a ground, but we see them in a particular grouping as well. Even simple patterns of For more Cengage Learning textbooks, visit www.cengagebrain.co.uk LOCALIZATION

get virtual fuel for their virtual jet by gripping and ungripping their healing hand ten times. Researchers at Shriners Childrens Hospital in Galveston (Flores et al., 2008) recently found that VR reduced pain during a passive range of motion exercises in children with large severe burn wounds. VR reduced patients pain for 25-minute physical therapy sessions, five days in a row, with no reduction in analgesic effectiveness. Three of the four pediatric burn patients showed large reductions in pain during VR, and one patient showed no reduction. Many patients report having fun during wound care and physical therapy, when allowed to use virtual reality. With funding from the Paul Allen Family Foundation, the National Institutes of Health, Scandinavian Design, and the Pfeiffer Foundation, Hoffman and worldbuilder Ari Hollander have developed a new more attention-grabbing virtual environment specifically designed for treating pain (selected into the 2006 Smithsonian Cooper-Hewitt National Museum of Design Triennial). Patients fly through an icy canyon with a river and frigid waterfall, and they shoot snowballs at snowmen, igloos, penguins and woolly mammoths (with animated impacts, sound effects, and soothing background music provided by Paul Simon). The technology for these advances in pain reduction are proceeding apace with the psychological advances. Hoffman, Jeff Magula and Eric Seibel have recently completed a custom optic fiber VR helmet that uses photons instead of electrons, so burn patients can get VR while sitting in the water-filled scrubtanks (Hoffman, Patterson et al., 2008). They also recently developed a pair of articulated robot-arm mounted helmet-less VR goggles for patients unable to wear conventional helmets (Maani, Hoffman et al., 2008). Hoffman, Patterson and colleagues are optimistic that virtual reality can provide a much needed psychological pain control technique that could prove valuable for treating other pain populations in addition to burn pain (e.g., combat-related blunt force trauma injuries, cancer procedures, emergency room 'ERVR', dental pain, and physical therapy during recovery from knee surgery). Their project nicely demonstrates the growing interdisciplinary alliance between research in psychology on the one hand and real-world problems in medicine on the other. Further details about the work can be found at www.vrpain.com.

dots fall into groups when we look at them. To illustrate this, begin by looking at the matrix of dots shown in Figure 5.5a. These dots are equally spaced up and down, so they can be seen as being organized in rows or columns, or even as lying along diagonal paths. This is, therefore, an ambiguous pattern that follows similar principles to those illustrated in Figures 5.3 and 5.4. Only one organization is seen at a time, and at intervals this organization will spontaneously switch to another. The Gestalt psychologists proposed a number of determinants of grouping for these kinds of dot patterns. For instance, if the vertical distance between dots is reduced, as in Figure 5.5b, columns will most likely be seen. This is grouping by proximity. If instead of varying the dot distances we vary the color shape of the elements,

160 CHAPTER 5 PERCEPTION Figure 5.4 The Slave Market with a Disappearing Bust of Voltaire. A reversible figure is in the center of this painting by Salvador Dali (1940). Two nuns standing in an archway reverse to form a bust of Voltaire. Salvador Dali, Slave Market with Disappearing Bust of Voltaire, 1950, The Salvador Dali Museum, St. Petersburg, Fla we can organize the dots on the basis of similarity (Figures 5.5c and d). If we move the dots to form two intersecting wave lines of dots, we are grouping by good continuation (Figure 5.5e), and if we enclose a space using lines of dots, we will tend to see grouping by closure. Note that in this last case we see a diamond positioned between two vertical lines, even though the pattern could be two familiar letters stacked on each other (W on M) or even facing each other (K and a mirror-image K). This illustrates the powerful nature of the Gestalt grouping determinants. These determinants serve to create the most stable, consistent, and simple forms possible within a given pattern. Modern research on

visual grouping has shown that the Gestalt determinants have a strong influence on perception. For example, in one series of studies, visual targets that were part of larger visual groupings based on proximity were much harder to detect than the same targets seen as standing outside the group (Banks & Prinzmetal, 1976; Prinzmetal, 1981). In another set of studies, targets that were dissimilar to nontargets in color and shape were easier to find than targets that were more similar (Triesman, 1986). Even the similarity among the various nontargets has an important effect: Targets are easier to find as the similarity of nontargets increases, allowing the target to 'pop out' as a figure distinct from the background (Duncan & Humphreys, 1989). Finally, there are reliable illusions associated with the Gestalt determinants, such that people judge distances among the elements within perceptual groups to be smaller than the same distances when they are between elements in different groups (Coren & Girgus, 1980; Enns & Girgus, 1985). All of these results show that visual grouping plays a large role in the way we organize our visual experience. For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

a) b) c) d) e) f) Figure 5.5 Gestalt Determinants of Grouping. (a) Equally spaced dots can be seen as rows, columns, or even diagonals. (b) Grouping into columns by proximity. (c) Grouping into columns by color similarity. (d) Grouping into columns by shape similarity. (e) Grouping by good continuation. (f) Grouping by closure. Although perceptual grouping has been studied mainly in visual perception, the same determinants of grouping appear in audition. Many demonstrations of this come from researchers who study music perception. Proximity in time clearly operates in audition. For example, four drumbeats with a pause between the second and third beats will be heard as two pairs. Similarly, sets of notes that are close together in time will be grouped together (as in the DUH-DUH-DUH-DUMMM opening of Beethoven's 5th symphony). Notes that are proximal in pitch will also be grouped together. Music often involves counterpoint, where two melodies occur simultaneously. Listeners can shift attention between melodies so that the attended

melody becomes the figure and the nonattended one becomes the ground. Often this is possible because the two melodies are in different octaves, so that notes within a melody are close to one another in pitch and notes between melodies are not. Similarity and closure are also known to play important roles in hearing tones and more complex stimuli (Bregman, 1990). Perceiving distance To know where an object is, we must know its distance or depth. Although perceiving an object's depth seems effortless, it is actually a remarkable achievement because we have no direct access to the depth dimension, thereby leading to one form of the many-to-one problem that we discussed earlier. A retina is a two-dimensional surface onto which a three-dimensional world is projected. The retina therefore directly reflects height and width, but depth information is lost and must somehow be reconstructed on the basis of subtle pieces of information known collectively as depth cues. Depth cues can be classified as binocular or monocular. Binocular cues Why are we and other animals equipped with two eyes rather than with just one? There are two reasons. Some animals, for example fishes, have eyes on either side of their head, which allows them to see a very large percentage of the world around them without moving their heads or their bodies. Other animals, for example humans, have two eyes in the front of their heads, both pointing in the same direction. Humans can see less of the world at any given instant than fishes, but they can use their two eyes to perceive depth. (Try covering one eye, and then sit as a passenger in a car driving in stop-and-go traffic. It's a scary experience, because you have much less sense than you normally would of how close you are to cars and other objects in front of you). The two eyes' ability to jointly infer depth comes about because the eyes are separated in the head, which means that each eye has a slightly different view of the same scene. You can easily demonstrate this by holding your right

index finger close to your face and examining it first with only one eye open and then with only the other eye open. The term binocular disparity is used to refer to the difference in the views seen by each eye. The disparity is largest for objects that are seen at close range and becomes smaller as the object recedes into the distance. Beyond 3–4 meters, the difference in the views seen by each eye is so small that binocular disparity loses its effectiveness as a cue for depth. However, for many everyday tasks, such as reaching for objects and navigating around obstacles, the difference in the views seen by each eye is a powerful cue for depth. In humans and other animals with binocular vision, the visual part of the brain uses binocular disparity to For more Cengage Learning textbooks, visit www.cengagebrain.co.uk LOCALIZATION assign objects to various locations in space, depending on how far apart the two images of an object are when compared. If the images of an object are in the same place in the two views, the brain assumes that this is the location on which both eyes are fixating. If the difference between the images is large, as it is for the two views of your finger held close to your face, the brain concludes that the object is much closer. In addition to helping us see depth in the everyday world, binocular disparity can be used to fool the eye into seeing depth when none is really present. One way this is achieved is by using a device called a stereoscope, which displays a slightly different photograph to each eye. In Victorian times these devices were proudly displayed in the sitting rooms of middle-class homes, much as highdefinition TV sets might be today. Yet the stereoscope is not just a curious antique. The same principle of binocular disparity is used today in children's 'View Master' toys, or in 'special effects' 3-D movies for which viewers must wear glasses with colored or light-polarizing filters that selectively allow one image to arrive at one eye and a slightly different image to arrive to the other. Monocular cues As indicated, the use of binocular cues is limited to objects that are relatively close. What about objects that are further away like distant clouds, cityscapes, or mountains? Here, binocular cues are relatively ineffective and other cues, known as monocular cues, must be used, and the task of the visual system is not straightforward. Essentially, the system has to make use of a hodge-podge of available information in the environment in order to come to a conclusion, much as a detective must use a hodge-podge of available evidence about a murder to figure out who the murderer is. Figure 5.7 illustrates five monocular cues; these plus one other are as follows.

1. Relative size. If an image contains an array of similar objects that differ in size, the viewer interprets the smaller objects as being farther away (see the trees in Figure 5.6).
2. Interposition. If one object is positioned so that it obstructs the view of the other, the viewer perceives the overlapping object as being nearer (see the buildings in Figure 5.6).
3. Relative height. Among similar objects, those that appear closer to the horizon are perceived as being farther away (see the birds in Figure 5.6).
4. Perspective. When parallel lines in a scene appear to converge in the image, they are perceived as vanishing in the distance (see the railroad tracks in Figure 5.6).
5. Shading and shadows. Whenever a surface in a scene is blocked from receiving direct light, a shadow is cast. If that shadow falls on a part of the same object that is

blocking the light, it is called an attached shadow or simply shading. If it falls on another surface that does not belong to the object casting the shadow, it is called a cast shadow. Both kinds of shadows are important cues to depth in the scene, giving us information about object shapes, distances between objects, and where the light source is in a scene (Coren, Ward, & Enns, 1999). 6. Motion. Have you ever noticed that if you are moving quickly – perhaps on a fast-moving train – nearby objects seem to move quickly in the opposite direction while more distant objects move

more slowly (though still in the opposite direction)? Extremely distant objects, such as the moon, appear not to move at all. The difference in the speeds with which these objects appear to move provides a cue to their distance from us and is termed motion parallax. Perceiving motion This last monocular cue, motion, brings us to the next main topic involving localization. If we are to move around our environment effectively, we need to know not only the locations of stationary objects but also the trajectories of moving ones. We need to know, for example, that the car coming toward us from a block away will not yet have arrived at the intersection by the time we have finished crossing the street. We must, that is, be able to perceive motion. Stroboscopic motion What causes us to perceive motion? The simplest idea is that we perceive an object is in motion whenever its image moves across our retina. This answer turns out to be too simple, though, for we can see motion even when nothing moves on our retina. This phenomenon, which is shown in Figure 5.7, was demonstrated in 1912 by Wertheimer in his studies of stroboscopic motion. Stroboscopic motion is produced most simply by flashing a light in darkness and then, a few milliseconds later, flashing another light near the location of the first light. The light will seem to move from one place to the other in a way that is indistinguishable from real motion. Wertheimer's demonstration of stroboscopic is not just an idle academic exercise; the phenomenon is crucial to great deal of present-day visual-display technology. A Height in field Interposition Shading/shadows Perspective Relative size Figure 5.6 Monocular Distance Cues in a Picture. Artists use some or all of these cues in combination to portray depth on a twodimensional surface. All of these cues are present in a photograph of a natural scene and are also present on the retinal image in the eye. © MACDUFF EVERTON/CORBIS CHAPTER 5 PERCEPTION For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

Percept Stimulus lights Time a) b) Figure 5.7 Stroboscopic Motion. The sequence of still frames in (a), shown at the appropriate intervals, results in the percept shown in (b). The illusion of continuous motion resulting from successively viewed still pictures is the basis of motion in movies, video, and television. prime example is movies, wherein the motion we perceive is stroboscopic motion. A movie is, as most people realize, simply a series of still photographs (or 'frames'), each one slightly different from the preceding one. Thus, as the frames are successively displayed on the screen, the discrete frame-to-frame differences in, say, the position of Daniel Craig's hand during an action sequence, in a James Bond film, are perceived as motion – stroboscopic motion to be sure, but motion which is perceived pretty much exactly as normal, continuous motion. a) This is an example of the types of displays used by investigators to study patterns of humans in motion. Positions of lights affixed to individuals are indicated. Figure 5.8 Patterns of Human Motion. For more Cengage Learning textbooks, visit www.cengagebrain.co.uk LOCALIZATION Real motion Of course, our visual system is also sensitive to real motion – that is, movement of an object through all intermediate points in space. However, the analysis of such motion under everyday conditions is amazingly complex. Some paths of motion on the retina must be attributed to movements of the eye over a stationary scene (as occurs when we are reading). Other motion paths must be attributed to moving objects (as when a bird enters our visual field). Moreover, some objects whose retinal images are stationary must be seen to be moving (as when we follow the flying bird with our eyes), while some objects whose retinal images are moving must be seen as stationary (as when the stationary background traces motion across the retina because our eyes are pursuing a flying bird). It therefore is not surprising that our analysis of motion is highly relative. We are much better at detecting motion when we can see an object against a structured background (relative motion) than when the background is a uniform color and only the moving object can be seen (absolute

motion). Certain patterns of relative movement can even serve as powerful cues to the shape and identity of three-dimensional objects. For example, researchers have found that the motion displays illustrated in Figure 5.8 are sufficient to enable viewers to easily identify the activity of a human figure, even though it consists of only 12 (or even fewer) points of light moving relative to one another (Johansson, von Hofsten, & Jansson, 1980). In Time b) A sequence of movement positions made

by a dancing couple.

164 CHAPTER 5 PERCEPTION CALIFORNIA MUSEUM OF PHOTOGRAPHY, UNIVERSITY OF CALIFORNIA, RIVERSIDE The Holmes-Bates stereoscope, invented by Oliver Wendell Holmes in 1861 and manufactured by Joseph Bates, creates a vivid perception of depth. other studies using these displays, viewers were able to identify their friends and even tell whether the model was male or female after seeing only the lights attached to the ankles (Cutting, 1986). Another important phenomenon in the study of real motion is selective adaptation. This is a loss in sensitivity to motion that occurs when we view motion; the adaptation is selective in that we lose sensitivity to the motion viewed and to similar motions, but not to motion that differs significantly in direction or speed. If we look at upward-moving stripes, for example, we lose sensitivity to upward motion, but our ability to see downward motion is not affected (Sekuler, 1975). As with other types of adaptation, we do not usually notice the loss of sensitivity, but we do notice the after-effect produced by adaptation. If we view a waterfall for a few minutes and then look at the cliff beside it, the cliff will appear to move upward. Most motions will produce such after-effects, always in the opposite direction from the original motion. How does the brain implement the perception of real motion? Some aspects of real motion are coded by specific cells in the visual cortex. These cells respond to some motions and not to others, and each cell responds best to one direction and speed of motion. The best evidence for the existence of such cells comes from studies with animals in which the experimenter records the responses of single cells in the visual cortex while the animal is shown stimuli with different patterns of motion. Such single-cell recording studies have found cortical cells that are tuned to particular directions of movement. There are even cells that are specifically tuned to detect an object moving toward the head, an ability that is clearly useful for survival (Regan, Beverly, & Cynader, 1979). These specialized motion cells provide a possible explanation for selective adaptation and the motion aftereffect. Presumably, selective adaptation to an upward motion, for example, occurs because the cortical cells that are specialized for upward motion have become fatigued. Because the cells that are specialized for downward motion are functioning as usual, they will dominate the processing and result in the aftereffect of downward motion. However, there is more to the neural basis of real motion than the activation of specific cells. We can see motion when we track a luminous object moving in darkness (such as an airplane at night). Because our eyes follow the object, the image is almost motionless on the retina, yet we perceive a smooth, continuous motion. Why? The answer seems to be that information about how our eyes are moving is sent from motor regions in the front of the brain to the visual cortex and influences the motion we see. In essence, the motor system is informing the visual system that it is responsible for the lack of regular motion on the retina, and the visual system then corrects for this lack. In more normal viewing situations, there are both eye movements and large retinal-image movements. The visual system must combine these two sources of information to determine the perceived motion. You can demonstrate a consequence of this arrangement ³ GINO SANTA MARIA j DREAMSTIME.COM

To control the ball and avoid being tackled, soccer players must be able to perceive motion accurately.

by gently pushing up on your eyeball through your lid. You'll note that the world appears to move. This is because the world is moving across your retina, but the normal signals from the motor regions are absent; the only way the brain can interpret it is if the world itself is moving.

INTERIM SUMMARY | To localize objects, we must first separate them and then organize them into groups. | Localization involves determining an object's position in the up-down and left-right dimensions. This is relatively easy because the required information is part of our retinal image. Localizing an object also requires that we know its distance from us. This form of perception, known as depth perception, is not so easy because it's not available in the retinal image. We have a variety of depth cues, both monocular and binocular, that allow us to do this. | Localizing an object sometimes requires that we know the direction in which an object is moving. This can be done either with real motion or with stroboscopic motion.

CRITICAL THINKING QUESTIONS 1 Imagine what your visual experience might be like if you suddenly became unable to see motion; in other words, suppose you saw things happening like a slide show than like a movie. How does motion perception contribute to your experience of a coherent world, and in what ways would the world become incoherent without a perception of motion. 2 Rank all the distance-perception cues from most important to least important. The main part of your answer should be to describe why you believe some distance-perception cues to be more or less important than others. This, of course, requires a definition on your part of what it means for a distance-perception cue to be 'important'.

RECOGNITION The perceptual system needs to determine not only where relevant objects are in the scene, but also what they are. This is the process of recognition. Ideally, if a cat crosses our path, we should be able to recognize it as a cat, not as a skunk or a hula hoop. Similarly, if a benign tent is in For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

RECOGNITION ^a ISTOCKPHOTO.COM/SO--COADDICT In the early stages of recognition, the perceptual system uses information on the retina to describe the object in terms of primitive components like lines and edges. In later stages, the system compares this description to those of various categories of objects stored in the visual memory, such as 'dogs'. front of us, we should be able to recognize it as a benign tent, not as a dangerous bear. (It is, however, noteworthy that from an evolutionary perspective, we would be better off misperceiving a tent for a bear than a bear for a tent. Our visual system has probably evolved in such a way that it is biased to perceive objects as dangerous even if sometimes they are not). Recognizing an object, in turn, entails several subproblems. First, we have to acquire fundamental or primitive features of information from the environment and assemble them properly. For example, if we acquire the information that there's something red and something green and a circle and a square we must somehow figure out that it's the circle that's red and the square that's green, not vice-versa. Second we have to figure out what the objects we're seeing actually are. In the simple example we've just described, we somehow have to figure out that it's a square there to begin with. A more complex

task would be to figure out that the combination of lines, angles, and shapes that we're looking at constitutes a human face, and a yet more complex task would be to figure out that the face belongs to a particular person, like Queen Elizabeth. In what follows, we will discuss these various functions of recognition. We'll start by talking about global-to-local processing: the means by which a scene aids in the perception of individual objects within the scene. We'll then move on to the binding problem: how activity in different parts of the brain, corresponding to different primitives

such as color and shape, are combined into a coherent perception of an object. Next, we'll talk about how we actually recognize what an object is. Global-to-local processing Look at the object in Figure 5.9 (left panel). What is it? It could be a loaf of bread or it could be a mailbox. How is the visual system to disambiguate these two possibilities? One of the most powerful tools used by the perceptual system to solve this and other similar problems is to use the context (the scene) within which the object is embedded to make inferences about what the object is. That is, the system can start by carrying out global processing – understanding what the scene is – followed by local processing – using knowledge about the scene to assist in identifying individual objects. Thus, if the system determined that the scene was of a street, the object would be interpreted as a mailbox, while if the system determined that the scene was of a kitchen, the object would be interpreted as a loaf of bread (see Figure 5.9a, middle and right panels). The logic of this process is articulated by Tom Sanocki (1993), who notes that objects in the world can appear in an infinitude of orientations, sizes, shapes, colors and so on and point out that, accordingly: 'If during object identification, the perceptual system considered such factors for an unconstrained set of alternatives, the enormous number of combinations of stimulus features and feature-object mappings would create a combinatorial explosion' (p. 878). Sanocki notes that an obvious means of reducing what would be an otherwise impossible information-processing task is to use early (global) information to constrain the interpretation of later information. A number of lines of research have determined that, indeed, exactly this kind of process occurs. For example, Schyns and Oliva (1994) showed composite pictures of naturalistic scenes. Composite pictures are 'double exposures' of two unrelated pictures, for example a skyline and a street. One of the scenes comprising the composite (say the skyline) contained only global information whereas the other (the street) contained only local information. These composites were then shown either briefly (e.g., around 10 milliseconds) or for longer (e.g., around 100 milliseconds) and the observers were asked what they had seen. For short exposures, observers reported seeing the scene containing only global information (the skyline in this example) while at the longer exposures, observers reported seeing the scene containing only local information (the street). This provides evidence that the visual system tends to acquire global information first, followed by local information. The binding problem: pre-attentive and attentive processes In our earlier discussion of attention, we learned that attention is the process by which we select which of the vast amount of incoming information is processed and eventually perceived consciously. Attention has also been conceptualized as having the role of binding together different features of an incoming stimulus. An excellent illustration of what we mean by this takes the form of what is known as an illusory conjunction. Suppose an observer is shown very briefly (e.g., a twentieth of a second) a stimulus such as the one in Figure 5.10 – a small red circle, a large green square, and a medium-size blue triangle – and asked to report what they saw. The observer is typically able to report the three shapes and the three colors – but often incorrectly reports which color went with which shape, e.g., the observer might

CHAPTER 5 PERCEPTION For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

report that the square was green, not red. Thus the conjunction of shape (square) and color (red) is what is perceived, but it is illusory. (People often experience a rough analogy of this phenomenon while reading: They might conjoin part of one word on one line of text, e.g., the 'liver' from 'delivery' with part of another word on a different line, e.g., the 'pool' from 'cesspool' and perceive that they see the word 'Liverpool' in the text – thereby misconjoining the primitive features of shape and location.) Feature integration theory Illusory conjunctions suggest that information from

the visual world is preattentively encoded along separate dimensions – in the example, shape and color are encoded separately – and then integrated in a subsequent attentive processing stage. This idea is, indeed, at the heart of feature-integration theory, initially proposed by Anne Treisman (Treisman, 1986, 1992). The general idea is that in a first, preattentive stage, primitive features such as shape and color are perceived while in the second, attentive stage, focused attention is used to properly ‘glue’ the features together into an integrated whole. Illusory conjunctions occur when stimulus duration is sufficient for the primitives to be obtained, but not sufficient for the longer, attentional gluing stage. A standard experimental procedure for distinguishing primitive features from ‘glued-together’ features is a visual search task in which the observer’s task is to determine whether some target object is present in a cluttered display. A typical visual search task is shown in Figure 5.11 where the task is to find a green ‘L’. In the left panel of Figure 5.11, the task is simple; the green L ‘pops out’ from the collection of red T’s and red L’s. In the right panel, however, the task of finding the same green L is considerably more difficult when the background is a collection of red L’s and green T’s. The reason, according to feature integration theory is that color is a primitive feature: In the left panel, you can simply scan the information all at once; what is red and what is green will perceptually separate and the presence of the one green object – the target green L – will be apparent. In the right panel, in contrast, you cannot distinguish the target from the background on the basis of the primitive attribute of color; you must attend to each letter, binding together the color and the shape, before you can determine whether that letter is or is not the target. Problems with feature integration theory Feature integration theory has enjoyed a great deal of support over the past couple of decades. In recent times, however, it has come under attack from the perspective of both theoretical parsimony and biological plausibility. The major problem is that, using visual search and related procedures, scientists have unveiled too many presumed ‘primitives’ to be realistic. A particularly lucid description of the problems with the theory is provided by Di Lollo, Kawahara, Suvic, and Visser (2001). They go on to describe an alternative, dynamic control theory whose central premise is that, ‘instead of an early, hard-wired system sensitive to a small number of visual primitives, there is a malleable system whose components can be quickly reconfigured to perform different tasks at different times, much as the internal pattern of connectivity in a computer is rearranged dynamically by enabling and disabling myriad gates under program control’ (p. 11). Figure 5.10 Illusory Conjunction. When images are flashed briefly, observers often miscombine shape and color. This is known as illusory conjunction. Figure 5.11 A Visual Search Task. Find the green L. This is any easy task in the left panel, where popout takes place, but a difficult task in the right panel, where each stimulus requires focal attention. RECOGNITION For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

168 CHAPTER 5 PERCEPTION This basically means that the system rearranges itself for different tasks – as opposed to there being many subsystems for each possible task. Determining what an object is Attentive versus preattentive processing is concerned with the problem of determining which visual characteristics belong to the same object. A second problem is that of using the resulting information to determine what an object actually is. Here, shape plays a critical role. We can recognize a cup, for example, regardless of whether it is large or small (a variation in size), brown or white (a variation in color), smooth or bumpy (a variation in texture), or presented upright or tilted slightly (a variation in orientation). In contrast, our ability to recognize a cup is strikingly affected by variations in shape; if part of the cup’s shape is hidden, we may not recognize it at all. One piece of evidence for the importance of shape is that we can recognize many objects about as well from simple line drawings, which preserve only the shapes of the objects, as well as from

detailed color photographs, which preserve many other attributes of the objects as well (Biederman & Ju, 1988). Here also, visual processing can be divided into earlier and later stages. In early stages, the perceptual system uses information on the retina, particularly variations in intensity, to describe the object in terms of primitive components like lines, edges, and angles. The system uses these components to construct a description of the object. In later stages, the system compares this description to those of various categories of objects stored in visual memory and selects the best match. To recognize a particular object as the letter B, for example, is to say that the object's shape matches that of B's better than it matches that of other letters.

Stimuli Feature detectors in the cortex Much of what is known about the primitive features of object perception comes from biological studies of other species (such as cats and monkeys) using single-cell recordings in the visual cortex. These studies examine the sensitivity of specific cortical neurons when different stimuli are presented to the regions of the retina associated with those neurons; such a retinal region is called a receptive field. These single-cell studies were pioneered by David Hubel and Torstein Wiesel (1968) who, in 1981, won a Nobel prize for their Nerve impulses For more Cengage Learning textbooks, visit www.cengagebrain.co.uk work. Hubel and Wiesel identified three types of cells in the visual cortex that can be distinguished by the features to which they respond. Simple cells respond when the eye is exposed to a line stimulus (such as a thin bar or straight edge between a dark and a light region) at a particular orientation and position within its receptive field. Figure 5.12 illustrates how a simple cell will respond to a vertical bar and to bars tilted away from the vertical. The largest response is obtained for a vertical bar, and the response decreases as the orientation varies from the optimal one. Other simple cells are tuned to other orientations and positions. A complex cell also responds to a bar or edge in a particular orientation, but it does not require that the stimulus be at a particular place within its receptive field. It will respond continuously as the stimulus is moved across that field. Hypercomplex cells require not only that the stimulus be in a particular orientation, but also that it be of a particular length. If a stimulus is extended beyond the optimal length, the response will decrease and may cease entirely. Since Hubel and Wiesel's initial reports, investigators have found cells that respond to shape features other than single bars and edges; for example, there are hypercomplex cells that respond to corners or angles of a specific length (DeValois & DeValois, 1980; Shapley & Lennie, 1985). All of the cells described above are referred to as feature detectors. Because the edges, bars, corners, and angles to which these detectors respond can be used to approximate many shapes, the feature detectors might be thought of as the building blocks of shape perception. As we will see later, though, this proposal seems to be more true of simple shapes like letters than of complex shapes like those of tables and tigers.

off on off on off on Figure 5.12 The Response of a Simple Cell. This figure illustrates the response of a simple cortical cell to a bar of light. The stimulus is on the top, the response on the bottom; each vertical spike on the bottom corresponds to one nerve impulse. When there is no stimulus, only an occasional impulse is recorded. When the stimulus is turned on, the cell may or may not respond, depending on the position and orientation of the light bar. For this cell, a horizontal bar produces no change in response, a bar at 45 degrees produces a small change, and a vertical bar produces a very large change.

Component Component Emergent Whole Feature A Feature B +

Property Object Figure 5.13 Relationships Between Features. When simple two-dimensional features such as lines, angles, and shapes are combined, the resulting pattern is highly dependent on the spatial relations between the component features. In addition, new features are created. These emergent features have a perceptual reality, even though they involve complex spatial relations. Relations among features There is more to a description of a shape than just its features: The relations among features must also be specified. The importance of such relations is illustrated in Figure 5.13 where it is evident that, for example, the features of a right angle and a diagonal line must be combined in a specific way to result in a triangle; likewise, a Y-intersection and a hexagon must be specifically aligned to result in the drawing of a cube. It was these kinds of relations between features that Gestalt psychologists had in mind when they emphasized that 'the whole is different from the sum of its parts'. One way in which the whole is different is that it creates new perceptual features that cannot be understood by simply examining the component parts. Figure 5.13 shows four such emergent features. These emerge from very specific spatial relations among more elementary features, but nevertheless often behave just like simpler features in perceptual tasks such as target detection and visual search (Enns & Resnick, 1990; Enns & Prinzmetal, For more Cengage Learning textbooks, visit www.cengagebrain.co.uk RECOGNITION 1984; He & Nakayama, 1992). These results indicate that the visual system performs many sophisticated analyses of shape before the results of these analyses are made available to consciousness. Later stages of recognition: network models Closure Now that we have some idea of how an object's shape is described, we can consider how that description is matched to shape descriptions stored in memory to find the best match - that is to decide what an object is. Horizontal surface Simple networks Much of the research on the matching stage has used simple patterns, specifically handwritten or printed letters or words. Figure 5.14 illustrates a proposal about how we store shape descriptions of letters. The basic idea is that letters are described in terms of certain features, and that knowledge about what features go with what letter is contained in a network of connections. Such proposals are referred to as connectionist models. These models are appealing in that it is easy to conceive how networks could be realized in the brain with its array of interconnected neurons. Thus, connectionism offers a bridge between psychological and biological models. The bottom level of the network in Figure 5.14 contains the features: ascending diagonal, descending diagonal, vertical line, and right-facing curve. The top level contains the letters themselves. We will refer to each of these features and letters as a node in the network. A connection between a feature and a letter node means that the feature is part of the letter. Connections ending in arrowheads are excitatory connections: If the feature is activated, the activation spreads to the letter (in a manner analogous to the way electrical impulses spread in a network of neurons).

Volume Convexity R P K Figure 5.14 A Simple Network. The bottom level of the network contains the features (ascending diagonal, descending diagonal, vertical line, and right-facing curve), the top level contains the letters, and a connection between a feature and a letter means that the feature is part of the letter. Because the connections are excitatory, when a feature is activated, the activation spreads to the letter.

170 CHAPTER 5 PERCEPTION To see how this network can be used to recognize (or match) a letter, consider what happens when the letter K is presented. It will activate the features of ascending diagonal, descending diagonal, and vertical line. All three of these features will activate the node for K, while two of them – the descending diagonal and vertical line – will activate the node for R; and one of them – the vertical line – will activate the node for P. Only the K node has all of its features activated, and consequently it will be selected as the best match. This model is too simple to account for many aspects of recognition, however. Consider what happens when the letter R is presented. It activates the features of descending diagonal, vertical line, and right-facing curve. Now the nodes for both R and P have all their features activated, and the model has no way of deciding which of the two categories provides a better match. What the model needs to know is that the presence of a descending diagonal means that the letter cannot be a P. This kind of negative knowledge is included in the augmented network in Figure 5.15, which has everything the preceding one had, plus inhibitory connections (symbolized by solid circles at their ends) between features and letters that do not contain those features. When a feature is connected to a letter by an inhibitory connection, activating the feature decreases activation of the letter. Thus, when R is presented to the network in Figure 5.16, the descending diagonal inhibits the P node, thereby decreasing its overall level of activation; now the R node will receive the most activation and, consequently, will be selected as the best match. Networks with feedback The basic idea behind the model we just considered – that a letter must be described by the features it lacks as well as by the features it contains – does not explain a pervasive and interesting finding: A letter is easier to perceive when it is presented as part of a word than when it is presented alone. For example, as shown in Figure 5.16, if observers are briefly presented with either the single letter K or the word WORK, they are more accurate in identifying whether a K or D was present R P K Figure 5.15 An Augmented Network. The network contains inhibitory connections between features and letters that do not contain these features, as well as excitatory connections. For more Cengage Learning textbooks, visit www.cengagebrain.co.uk Masking field and response alternatives D K Word or letter Fixation point WORK Time Figure 5.16 Perception of Letters and Words. This figure illustrates the sequence of events in an experiment that compares the perceptibility of a letter presented alone or in the context of a word. First, participants saw a fixation point, followed by a word or a single letter, which was present for only a few milliseconds. Then the experimenter presented a stimulus that contained a visual mask in the positions where the letters had been, plus two response alternatives. The task was to decide which of the two alternatives occurred in the word or letter presented earlier. (After Reicher, 1969) when the display contained a word than when it contained only a letter. To account for this result, our network of feature-letter connections has to be altered in a few ways. First, we have to add a level of words to our network, and along with it excitatory and inhibitory connections that go from letters to words, as shown in Figure 5.17. In addition, we have to add excitatory connections that go from words down to letters; these top-down feedback connections explain why a letter is more perceptible when presented briefly in a word than when presented briefly alone. When R is presented alone, for example, the features of vertical line, descending diagonal, and right-facing curve are activated, and this activation spreads

to the node for R. Because the letter was presented very briefly, not all the features may have been fully activated, and the activation culminating at the R node may not be sufficient for recognition to occur. In contrast, when R is presented in RED, there is activation not only from the features of R to the R node, but also from the features of E and D to their nodes; all of these partially activated letters then partially activate the RED node, which in turn feeds back activation to its letters via its top-down connections. The upshot is that there is an additional source of activation for R when it is presented in a word – namely, activation coming from the word – and this is why it is

RED PET R P K Figure 5.17 A Network with Top-Down Activation. The network contains excitatory and inhibitory connections between letters and words (as well as between features and letters), and some of the excitatory connections go from words to letters. easier to recognize a letter in a word than when it is presented alone. Many other findings about letter and word patterns have been shown to be consistent with this connectionist model (McClelland & Rumelhart, 1981). Models like these have also been used successfully in machines designed to read handwriting and recognize speech (Coren, Ward, & Enns, 1999). Recognizing natural objects and top-down processing a) Geons We know quite a bit about the recognition of letters and words, but what about more natural objects – animals, plants, people, furniture, and clothing? In this section we examine how we recognize such objects. Features of natural objects The shape features of natural objects are more complex than lines and curves, and more like simple geometric forms. These features must be such that they can combine to form the shape of any recognizable object (just as lines and curves can combine to form any letter). The features of objects must also be such that they can be determined or constructed from more primitive features, such as lines and curves, because, as noted earlier, primitive features are the only information available to the system in the early stages of recognition. For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

RECOGNITION One popular though controversial suggestion is that the features of objects include a number of geometric forms, such as cylinders, cones, blocks, and wedges, as illustrated in Figure 5.18a. These features, referred to as geons (short for ‘geometric ions’), were identified by Biederman (1987), who argues that a set of 36 geons, such as those in Figure 5.18a, combined according to a small set of spatial relations, is sufficient to describe the shapes of all objects that people can possibly recognize. To appreciate this point, as shown in Figure 5.18b you can form an object by combining any two geons – and the number of possible such two-geon objects is $36 \times 36 \times \frac{1}{4} = 1,296$; likewise, the number of possible three-geon objects is $36 \times 36 \times 36 \times \frac{1}{4} = 46,656$. Thus, two or three geons are sufficient to create almost 50,000 objects, and we have yet to consider objects made up of four or more geons. Moreover, geons like those in Figure 5.18a can be distinguished solely in terms of primitive features. For example, geon 2 in Figure 5.18a (the cube) differs from geon 3 (the cylinder) in that the cube has straight edges but the cylinder has curved edges; straight and curved edges are primitive features. Evidence that geons are features comes from experiments in which observers try to recognize briefly presented objects. The general finding is that recognition of an object is good to the extent that the geons of the object are perceptible. In one study, part of the shape of an object was deleted in such a way that the deletion either interfered with recovering the geons (see the right column of Figure 5.19) or did not (see the middle column). Recognition of the objects was much better when there was no interference with the geons. b) Objects 5 2 3 3 5 3 2 4 3 When the features (geons) are combined, they form natural objects. Note that when the arc (geon 5) is connected to the side of the cylinder (geon 3), it forms a cup; when connected to the top of the cylinder, it forms a pail. Wedges, cubes, cylinders, cones,

and arcs may be features of complex objects. Figure 5.18 A Possible Set of Features (Geons) for Natural Objects. (From L. Biederman, *Computer Vision, Graphics, and Image Processing*, 32, pp. 29-73, © 1985 Academic Press. Used with permission.)

172 CHAPTER 5 PERCEPTION Figure 5.19 Object Recognition and Geon Recovery. Items used in experiments on object recognition. The left column shows the original intact versions of the objects. The middle column shows versions of the objects in which regions have been deleted, but the geons are still recoverable. The right column shows versions of the objects in which regions have been deleted and the geons are not recoverable. Recognition is better for the middle versions than for the rightmost versions. (After L. Biederman, *Computer Vision, Graphics, and Image Processing*, 32, pp. 29-73, © 1985 Academic Press. Used with permission.) As usual, the description of an object includes not just its features but also the relations among them. This is evident in Figure 5.18b. When the arc is connected to the side of the cylinder, it forms a cup; when it is connected to the top of the cylinder, it forms a pail. Once the description of an object's shape is constructed, it is compared to an array of geon descriptions stored in memory to find the best match. This matching process between the description of an object's shape and the descriptions stored in memory resembles the process described earlier for letters and words (Hummel & Biederman, 1992). The importance of context A key distinction in perception, to which we have previously alluded, is that between bottom-up and top-down processes. Bottom-up processes are driven solely by the input – the raw, sensory data – whereas top-down processes are driven by a person's knowledge, experience, For more Cengage Learning textbooks, visit www.cengagebrain.co.uk attention, and expectations. To illustrate, recognizing the shape of an object solely on the basis of its geon description involves only bottom-up processes; one starts with primitive features of the input, determines the geon configuration of the input, and then makes this description available to shape descriptions stored in memory. In contrast, recognizing that the object is a lamp partly on the basis of its being on a night table next to a bed involves some top-down processes; other information is used besides the input regarding shape. While most of the processes considered thus far in this chapter are bottomup ones, top-down processes also play a major role in object perception. Top-down processes, in the form of expectations, underlie the powerful effects of context on our perception of objects and people. You expect to see your chemistry lab partner, Sarah, every Tuesday at 3 p.m., and when she enters the lab at that moment you hardly need to look to know it is she. Your prior knowledge has led to a powerful expectation, and little input is needed for recognition. But should Sarah suddenly appear in your hometown during Christmas vacation, you may have trouble recognizing her. She is out of context – your expectations have been violated, and you must resort to extensive bottom-up processing to tell that it is in fact she (we experience this as 'doing a double take'). As this example makes clear, when the context is appropriate (that is, it predicts the input object), it facilitates perception; when the context is inappropriate, it impairs perception. Experimental evidence for the role of context in object perception comes from semantic priming studies. Here a to-be-identified stimulus (e.g., the word DOCTOR) is briefly preceded by a priming stimulus that is either related to it (e.g., NURSE) or unrelated (e.g., CHAIR); studies have shown that both pictures and words are identified more quickly and remembered more accurately when they are preceded by related rather than unrelated primes (e.g., Palmer, 1975; Reinitz, Wright, and Loftus, 1989). The effects of context are particularly striking when the stimulus object is ambiguous – that is, can be perceived in more than one way. An ambiguous figure is presented in Figure 5.20; it can be perceived either as an old woman or as a young woman. If you have been looking at unambiguous pictures that resemble the

young woman in the figure (that is, if young women are the context), you will tend to see the young woman first in the ambiguous picture. This effect of temporal context is illustrated with another set of pictures in Figure 5.21. Look at the pictures as you would at a comic strip, from left to right and top to bottom. The pictures in the middle of the series are ambiguous. If you view the figures in the sequence just suggested, you will tend to see the ambiguous pictures as a man's face. If you view the figures in the opposite order, you will tend to see the ambiguous pictures as a young woman.

Context effects and top-down processing also occur with letters and words, and play a major role in reading. Both the number of eye fixations we make on text and the durations of these fixations are greatly influenced by how much we know about the text – and, hence, by the amount of top-down processing we can invoke. When the material is unfamiliar, there is little top-down processing. In such cases we tend to fixate on every word, except for function words like 'a', 'of', 'the', and so on. As the material becomes more familiar, we can bring our prior knowledge to bear on it, and our fixations become shorter and more widely spaced (Just & Carpenter, 1980; Rayner, 1978). Top-down processing occurs even in the absence of context if the input is sufficiently sparse or degraded. Suppose that at a friend's apartment you enter her dark kitchen and see a smallish black object in the corner. You think the object could be your friend's cat, but the perceptual input is too degraded to convince you of this, so you think of a particular feature of the cat, such as its tail, and selectively attend to the region of the object that is likely to contain that feature if it is indeed a cat (Kosslyn & Koenig, 1992). This processing is top-down, because you have used specific knowledge – the fact that cats have tails – to generate an expectation, which is then combined with the visual input. Situations like this are common in everyday life. Sometimes, however, the input is very degraded and the expectations we form are way off the mark, as when we finally realize that our would-be cat in the kitchen is really our friend's purse. As the previous discussion makes clear, one reason that top-down processing is useful is because it constrains the

Figure 5.21 Effects of Temporal Context. What you see here depends on the order in which you view the pictures. If you start at the beginning and work forward, the middle pictures will appear to be a young woman. In other words, your initial perception perseverates. (From G. H. Fisher (1967) 'Perception of Ambiguous Stimulus Materials', from *Perception & Psychophysics*, 2:421–422. Reprinted by permission of the Psychonomic Society.)

Figure 5.20 An Ambiguous Stimulus. An ambiguous drawing that can be seen either as a young woman or as an old woman. Most people see the old woman first. The young woman is turning away, and we see the left side of her face. Her chin is the old woman's nose, and her necklace is the old woman's mouth. RECOGNITION For more Cengage Learning textbooks, visit www.cengagebrain.co.uk From *American Journal of Psychology*. Copyright 1930 by the Board of Trustees of the University of Illinois. Used with permission of the author and the University of Illinois Press

174 CHAPTER 5 PERCEPTION set of objects that are likely to occur in a given setting. For instance, we do not mistake a loaf of bread in a kitchen as a mailbox because we know that bread and not mailboxes tend to occur in kitchens. Similarly, individuals more accurately identify spoken words when they can see the speaker's lips than when they can't see them, because we have learned that specific lip movements constrain the set of sounds that the speaker can produce (e.g., Sams et al., 1991). However, these same top-down processes can sometimes produce perceptual illusions such that our perceptions are distorted by our expectations. One interesting example, called the McGurk effect (McGurk & MacDonald, 1976) results from conflicting auditory and visual

information. In particular, an observer watches a video of a speaker in which the a speaker's lips form the sound, 'ga-ga', while the simultaneous sound track provides speech that is normally perceived as 'ba-ba'. These sources of information are in conflict because we have learned that it is not possible to produce the sound 'ba' without closing one's lips; however, because the video portrays the speaker mouthing 'ga' his lips never close. The conjunction of these conflicting sources of information, surprisingly, produces the perception of 'da-da'. Thus, the speaker integrates the visual and auditory information with an entirely unexpected, 'illusory' result. Perceptual distortions resulting from top-down processes may sometimes lead to tragedy. In 1996 New York City police chased an African man named Amadou Diallo to his doorway. Thinking that the police were asking him for identification he took his wallet from his pocket and was instantly killed in a barrage of bullets from police officers who apparently thought that he had drawn a gun. Motivated by this and similar tragedies psychologists have developed video-game-like procedures to investigate such misperceptions. In a typical experiment people are told to shoot individuals on the screen who draw guns, but not individuals who brandish harmless objects. Studies have repeatedly shown that stereotypes strongly influence performance in this task; participants are more likely to shoot people with dark skin than to shoot light-skinned people when they quickly draw harmless objects (e.g., Correll et al., 2002, Dasgupta, McGhee, Greenwald, & Banaji, 2000). These simulations have been helpful in training police officers to avoid these potential biases.

Special processing of socially relevant stimuli: face recognition As the Diallo case demonstrates, social factors can influence perception. In fact, evidence suggests that people have developed perceptual processes that are specialized for processing socially relevant stimuli. Nowhere is this more true than in recognizing faces. It is of the utmost For more Cengage Learning textbooks, visit www.cengagebrain.co.uk social importance to be able to recognize kin, and to distinguish friend from foe. In addition, faces tend to be similar to one another. While other types of objects, such as houses, can differ in terms of the number and location of features (e.g., houses can have doors and windows in diverse places) faces all contain eyes, a nose, and a mouth in the same general pattern. The social importance of faces, combined with inherent recognition difficulties resulting from their similarity to one another, has apparently led to the development of special recognition processes that are employed for faces but not for objects. Three types of evidence are often cited as evidence for special face processing. First, prosopagnosia is a syndrome that can arise following brain injury, in which a person is completely unable to identify faces but retains the ability to recognize objects. Second, the inversion effect (Yin, 1969, 1970) is the name given to the finding that faces but not objects are extremely hard to recognize when they are presented upside-down, such as the photograph below. Finally, object recognition and face recognition appear to

³ MARTIN JENKINSON/ALAMY The former UK Prime Minister Tony Blair

have different developmental trajectories. Childrens' abilities to recognize objects tends to increase steadily with age; however, there is evidence that for many children face recognition ability actually declines temporarily during early adolescence. A popular theory to account for these face-object differences is that while objects are recognized on the basis of their component parts, faces are recognized on the basis of the overall pattern (or configuration) that the parts form (e.g., Farah, Tanaka, & Drain, 1995). By this explanation, prosopagnosics retain the ability to perceptually process parts but not configurations (e.g., Sergent, 1984), and inversion obscures parts less than it obscures the overall pattern that the parts form (Rock, 1988). Failure of recognition Recognizing an object is usually so automatic and effortless that we take it for granted. But the process sometimes breaks down. We have already seen that in normal people, recognition can fail in simple situations

(as with illusory conjunctions) and in more complex situations (as when a tent is mistaken for a bear). Recognition also fails routinely in people who have suffered from certain kinds of brain damage (due to accidents or diseases such as strokes). The general term for such breakdowns or disorders in recognition is agnosia. Of particular interest is a type of agnosia called associative agnosia. This is a syndrome in which patients with damage to temporal lobe regions of the cortex have difficulty recognizing objects only when they are presented visually. For example, the patient may be unable to name a comb when presented with a picture of it, but can name it when allowed to touch it. The deficit is exemplified by the following case. For the first three weeks in the hospital the patient could not identify common objects presented visually and did not know what was on his plate until he tasted it. He identified objects immediately on touching them [but] when shown a stethoscope, he described it as 'a long cord with a round thing at the end', and asked if it could be a watch. He identified a can opener as 'could be a key'. Asked to name a cigarette lighter, he said, 'I don't know'. He said he was 'not sure' when shown a toothbrush. Asked to identify a comb, he said, 'I don't know'. For a pipe, he said, 'some type of utensil, I'm not sure'. Shown a key, he said, 'I don't know what that is; perhaps a file or a tool of some sort'. (Reubens & Benson, 1971) What aspects of object recognition have broken down in associative agnosia? Since these patients often do well on visual tasks other than recognition – such as drawing For more Cengage Learning textbooks, visit www.cengagebrain.co.uk RECOGNITION objects or determining whether two pictured objects match – the breakdown is likely to be in the later stages of recognition, in which the input object is matched to stored object descriptions. One possibility is that the stored object descriptions have been lost or obscured in some way (Damasio, 1985). Some patients with associative agnosia have problems recognizing certain categories but not others. These category-specific deficits are of considerable interest because they may tell us something new about how normal recognition works. The most frequent category-specific deficit is loss of the ability to recognize faces, called prosopagnosia. (We discussed this condition briefly in Chapter 1.) When this deficit occurs, there is always brain damage in the right hemisphere and often some damage in homologous regions of the left hemisphere as well. The condition is illustrated by the following case. He could not identify his medical attendants. 'You must be a doctor because of your white coat, but I don't know which one you are. I'll know if you speak'. He failed to identify his wife during visiting hours... . He failed to identify pictures of Churchill, Hitler, and Marilyn Monroe. When confronted with such portraits he would proceed deductively, searching for the 'critical' detail which would yield the answer. (Pallis, 1955) A second kind of category deficit is loss of the ability to recognize words, called pure alexia (typically accompanied by damage in the left occipital lobe). Patients with this deficit typically have no difficulty recognizing natural objects or faces. They can even identify individual letters. What they cannot do is recognize visually presented words. When presented with a word, they attempt to read it letter by letter. It can take as long as ten seconds for them to recognize a common word, with the amount of time needed increasing with the number of letters in the word (Bub, Blacks, & Howell 1989). Other types of category-specific deficits involve impairment in the ability to recognize living things such as animals, plants, and foods. In rare cases patients are unable to recognize nonliving things such as household tools (Warrington & Shallice, 1984). Some of the suggested explanations of category-specific deficits have implications for normal recognition. One hypothesis is that the normal recognition system is organized around different classes of objects – one subsystem for faces, another for words, a third for animals, and so on – and these subsystems are localized in different regions of the brain. If a patient suffers only restricted brain damage, he or she may show a loss of one subsystem but not others. Damage in a specific part of the right hemisphere, for example, might disrupt the facerecognition subsystem but

leave the other subsystems intact (Damasio, 1990; Farah, 1990).

176 CHAPTER 5 PERCEPTION INTERIM SUMMARY | Recognizing an object requires that the various features associated with the object (such as shapes and colors) be correctly bound together, a process that requires attention. | Recognition of a particular object is aided by first acquiring 'global' aspects of the scene; for example quickly understanding that you are looking at a kitchen helps recognizing an ambiguous object as a loaf of bread rather than a mailbox. | Recognizing an object entails binding together various features of an object such as its shape and its color. The features themselves are acquired via pre-attentive processes, while 'gluing' them together require attention. | There are known kinds of cells in the visual cortex that are sensitive to various kinds of stimulus features such as orientation and position within the visual field. | Recognition of visual stimuli can be mimicked by a connectionist model or network. | Bottom-up recognition processes are driven solely by the input, whereas top-down recognition processes are driven by a person's knowledge and expectations. The shape features of natural objects are more complex than lines; they are similar to simple geometric forms such as cylinders, cones, blocks, and wedges. A particular set of such forms is a geon. | Face recognition may be special, i.e., different in important respects from recognition of other objects. CRITICAL THINKING QUESTIONS 1 At the beginning of this chapter we described a tent that was tragically mistaken for a bear. Why do you think this misperception happened. What could the hunters have done to have avoided the misperception? 2 Do you think there is a fundamental difference between recognizing a natural object, such as an eagle, and recognizing an artificial object such as a stop sign? Give reasons for your answers. ABSTRACTION The physical description of an object is a listing of all the information necessary to completely reproduce the object. Many stimuli studied in the scientific laboratory - patches of light, squares, single letters - are relatively simple and For more Cengage Learning textbooks, visit www.cengagebrain.co.uk ^a KURT KRIEGER/CORBIS Australian actress Cate Blanchett their physical descriptions are likewise simple. However, the physical description of most real-life, natural objects is enormously complicated. Look at Cate Blanchett pictured above. The visual detail that exists within it seems almost infinite. As you look closer and closer at her skin, for example, small blemishes and irregularities become apparent. Each individual hair on her head is positioned just so. The shadowing across her features, while subtle, is complex. To write a complete description of her face, in other words, would take an extremely long time. Really the only way you could do it would be by creating a bitmap of her face, and even then, the completeness of the description would be limited by the bitmap's resolution. Exact to abstract However, in real life, these limitations don't usually present a problem because you don't need all that much detail to solve the problems assigned to you by the world. For instance, in the Cate Blanchett example, you would only need as much detail as is necessary to (1) recognize her face to begin with and (2) determine from her expression what kind of mood she is in (caricaturists

know this quite well; with a few deft strokes of their pen, they can capture the likeness and expression of a person with remarkable clarity). This situation is not, of course, unique to faces. Whether you are looking at a hairpin or a pencil sharpener or an armchair or anything else, you rarely if ever need to know all the infinite visual detail. Rather, you only need to know enough to carry out whatever task is requiring you to perceive the object to begin with. The advantages of abstraction: required storage and processing speed To get a feel for this, look at the two drawings in Figure 5.22. Both were created using a computer drawing program. The face on the left was drawn freehand, while the one on the right was drawn as a 'copy' of the one on the left, using

nothing but the drawing program's oval and line tools. Clearly the left-hand original contains considerably more detail; however both give the same impression – of a slightly bewildered looking individual. When these two versions of the face were saved as files, the original, freehand version required 30,720 bytes of memory, while the 'abstracted' version required only 902 bytes – a savings of about 97 percent! Clearly it is more efficient in many respects to perceive and encode in memory an abstraction of the object rather than an exact representation of the object itself. As we noted earlier, object recognition is well conceptualized as the construction of objects using a 'drawing program' where the primitives are geons. A nice example of how perception of a real-life object is schematized in this manner was reported by Carmichael, Hogan, & Walter (1932), who presented ambiguous stimuli such as those shown in Figure 5.23. middle column, labeled 'Stimulus Figures', along with a label that told the observers what they were looking at. For instance, while viewing the stimulus in the top middle column, some observers were told that they were looking at 'curtains in a window' while others were told that were looking at 'a diamond in a rectangle'. The observers were later asked to reproduce what they had seen. Examples are shown in the left and right columns of Figure 5.23. As you can see, what the subjects perceived and stored in memory corresponded very strongly to what they considered themselves to be looking at. A more recent, and quite different demonstration of abstraction was reported by Intraub and Richardson (1989). Here, observers were shown pictures of objects such as those shown in the top panels of Figure 5.24. the general finding was that when the observers later redrew the pictures, they expanded the boundaries, as shown in the bottom panels of Figure 5.24. The conclusion again is that, rather than perceiving, storing, and later remembering a more-or-less literal image of what they had seen, the observers abstracted the important information (here the object's context as well as the object itself). The notion of abstraction harks back to our discussion in Chapter 4 of color metamers. You'll recall that color metamers are different physical stimuli (for instance a pure yellow light on the one hand and a redgreen mixture on the other) that lead to the exact same color perception. In this instance, the visual system is throwing away the information corresponding to the physical difference between the stimuli. Abstraction entails much the same thing: The information corresponding to the exact physical description (the 'bitmap') of the stimulus is lost; what is retained is the critical information that is needed. Figure 5.22 The Process of Abstraction. Two versions of the same sad face. The one on the left was drawn freehand, and the one on the right was drawn with 'abstracting' tools such as ovals and lines. The left face takes up considerably more disk space than the right, which illustrates one of the virtues of abstracting for any visual-processing device, including biological visual systems. ABSTRACTION For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

INTERIM SUMMARY | Abstraction is the process of converting the raw sensory information acquired by the sense organs (for example, patterns of straight and curved lines) into abstract categories that are pre-stored in memory (for example, letters or words). | Abstracted information takes less space and is therefore faster to work with than raw information. A useful analogy is between a bitmapped computer image of a face versus an abstracted image of the same face that is made up of preformed structures such as ovals and lines. CRITICAL THINKING QUESTIONS 1 In what way is the behavior of a visual artist influenced by color and shape constancy? Can you think of ways in which perceptual constancies actually make the artist's task more difficult than it would be without constancy? 2 In Chapter 4 we talked about metamers. Can you see a relation between metamers and the process of abstraction? What is it? Curtains in a window Diamond in a rectangle Four Sun Table Canoe Trowel Broom Eight Seven Gun Two Hourglass Kidney bean Pine tree Ship's wheel

Reproduced figure Verbal labels Stimulus figures Verbal labels Reproduced figure Figure 5.23 Verbal Labels and Abstraction. Carmichael, Hogan, and Walter (1932) showed people the kind of ambiguous stimuli shown in the middle panel. Observers were given one of the two verbal labels shown in the second and fourth columns. The subjects' later reconstructions of what they had seen conformed to the verbal label, as shown in the first and fifth columns. This experiment indicates that subjects remember not what they literally saw but rather abstract the fundamental information from it. CHAPTER 5 PERCEPTION For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

PERCEPTUAL CONSTANCIES You walk into a movie and discover, somewhat to your annoyance, that because all the seats in the middle section of the theater are taken, you are forced to sit far over on the left side. As the movie begins, however, you forget about your seating locale and just lose yourself in the movie's plot, its characters and its stunning special effects. All visual aspects of the movie appear to be entirely normal – and yet they're not. Because you're sitting off to the side, at an angle to the screen, the image of the movie screen on your retina is not a rectangle; rather it's a trapezoid, and all the visual images you see on the screen are analogously distorted. And yet this doesn't really bother you; you see everything as normal. How can this be? In this section we will describe a truly remarkable ability of the perceptual systems, termed the maintenance of constancy. The nature of constancies To understand the idea of constancies, it is important to first understand the relation and distinction between the inherent physical characteristics of an object and the information available to our perceptual systems about these objects. A movie screen, for example, is rectangular; that's a physical characteristic of it. But the image of it on our retina can be rectangular or trapezoidal depending on the angle from which you view it. A black cat seen in bright light is objectively lighter (it reflects more light to you) than a white cat in dim light; yet somehow in any kind of light, we maintain the perception that the black cat is actually black, while the white cat is actually white. An elephant seen from far away projects a smaller image on our retina than a gopher seen from close up; yet somehow, no matter what the distance, we maintain the perception that the elephant is larger than the gopher. In general, what we perceive is – and this almost sounds like magic – a perception of what an object is actually like rather than a perception based solely on the 'objective' physical information that arrives from the environment. Although constancy is not perfect, it is a salient aspect of visual experience and it should be; otherwise the world would be one where sometimes elephants are smaller than mice and where Denzel Washington is sometimes lighter Figure 5.24 Boundary Extension and Abstraction. Subjects tend to remember having seen a greater expanse of a scene than was shown to them in a photograph. For example, when drawing the close-up view in panel A from memory, the subject's drawing (panel C) contained extended boundaries. Another subject, shown a wider-angle view of the same scene (panel B), also drew the scene with extended boundaries (panel D). PERCEPTUAL CONSTANCIES For more Cengage Learning textbooks, visit www.cengagebrain.co.uk Reprinted by permission of Intraub and Richardson (1989), © Journal of Experimental Psychology: Learning, Memory, and Cognition

180 CHAPTER 5 PERCEPTION ^a DAVID FRAZIER/THE IMAGE WORKS Perceptual constancy enables us to determine how far away objects are. colored than Brad Pitt, depending on the particular situation. If the shape and color of an object changed every time either we or it moved, the description of the object that we construct in the early stages of recognition would also change, and recognition would become an impossible task. Color and brightness constancy Suppose I tell

you that I am thinking of two numbers whose product is 36, and I ask you to tell me what the two original numbers are. Your reasonable response would be that you don't have enough information to answer: The numbers I'm thinking of could be 2 and 18, or 6 and 6, or any of an infinite number of other pairs. Impossible though this task seems, it is, in a very real sense, what the visual system does when it maintains lightness and color constancy. To see what we mean by this, suppose you are looking at something, say a piece of red paper, and asked to name its color. Color constancy refers to the fact that you would report the paper to be red whether it were inside a room lit by an incandescent bulb, For more Cengage Learning textbooks, visit www.cengagebrain.co.uk which illuminates the paper with one particular set of wavelengths or outside in the noonday sun, which illuminates the paper with a very different set of wavelengths. It stands to reason that the perceived redness of the red paper is based on the wavelengths of the light that is reflected off the paper reaching your eyes. We will call these the available wavelengths. Let's now consider the physics of where these available wavelengths come from. It's a two-step process. First, the paper is illuminated by some light source which could be, among many other things, an incandescent bulb inside, or the sun outside. We will call the wavelengths provided by the source the source wavelengths. Second, the red paper itself reflects some wavelengths more than others (in particular it reflects mostly wavelengths corresponding to red and less of other wavelengths). We will call this property of the paper the reflectance characteristic. Now in a very real, mathematical sense, the available wavelengths reaching your eyes is the product of the source wavelengths and the reflectance characteristic. Realizing this puts us in a position to define color constancy, which is the ability of the visual system to perceive the reflectance characteristic – an inherent property of the object – no matter what the source wavelengths. It is in this sense, therefore that the visual system is presented with a product – the available wavelengths – and somehow figures out one of the factors, namely the reflectance characteristic. The incandescent bulb and the sun provide very different source wavelengths and – because the reflectance characteristic of the red paper doesn't change – very different available wavelengths therefore reach the eye. Yet somehow, the visual system is able to divide the source wavelengths out of the available wavelengths to arrive at the correct reflectance characteristic in both cases. This is analogous to your somehow figuring out that the first number I'm thinking of (analogous to the source wavelengths) is 12 which means that the other number (analogous to the reflectance characteristic) must be $36 / 12$ or 3. Brightness constancy is similar to color constancy, and refers to the fact that the perceived lightness of a particular object changes very little, if at all, even when the intensity of the source, and thus the amount of light reflected off the object changes dramatically. Thus, a black velvet shirt can look just as black in sunlight as in shadow, even though it reflects thousands of times more light when it is directly illuminated by the sun. A dramatic example of this finding is shown in the left-hand checkerboard picture on the next page: The squares labeled A and B are, astonishingly, exactly the same level of grey. We have demonstrated this in the right-hand version which is identical except that the two squares have been connected by gray bars. Your visual system is responding though, not to the physical data arriving at your eyes, but rather to the data plus the visual system's inferences about the grey level of the square: it 'corrects'

for the shadow being cast on Square B with a resulting perception of a white square that is as white as any of the other white portions of the board! How does the visual system manage to do these tricks? A clue comes about by examining the circumstances under which constancy fails. Suppose that the black shirt is put behind an opaque black screen and you view the shirt through a peephole in the screen. The screen reduces what you see through the opening to just the actual

light reflected from the shirt, independent of its surroundings. Now, when it is illuminated, the shirt looks white because the light that reaches your eye through the hole is more intense than the light from the screen itself. This demonstration underscores the fact that when we perceive objects in natural settings, rather than through peepholes, many other objects are usually visible. Color and brightness constancy depend on the relations among the intensities of light reflected from the different objects; essentially by using our past knowledge of object colors in general, our visual system is able to correct for the effect of the source illumination (both the source intensity and the source wavelengths) and arrive at the brightness and the color of the objects being seen (Gilchrist, 1988; Land, 1977; Maloney & Wandell, 1986).

Shape constancy We have provided an example of shape constancy in describing the non-effect of sitting to one side of a movie theater. Another is illustrated in Figure 5.25. When a door swings toward us, the shape of its image on the retina goes through a series of changes. The door's rectangular shape produces a trapezoidal image, with the edge toward us wider than the hinged edge; then the trapezoid grows thinner, until finally all that is projected on the retina is a vertical bar the thickness of the door. Nevertheless, we perceive an unchanging door swinging open. The fact that the perceived shape is constant while the retinal image changes is an example of shape constancy.

Size constancy The most thoroughly studied of all the perceptual constancies is size constancy: An object's perceived size remains relatively constant no matter how far away it is. As an object moves farther away from us, we generally do not see it as decreasing in size. Hold a quarter 1 foot in front of you and then move it out to arm's length. Does it appear to get smaller? Not noticeably. Yet, as shown in Figure 5.26, the retinal image of the quarter when it is 24 inches away is only about half the size of its retinal image when it is 12 inches away.

Dependence on depth cues The example of the moving quarter indicates that when we perceive the size of an object, we consider something in addition to the size of the retinal image. That additional something is the perceived distance of the object. As long ago as 1881, Emmert was able to show that size

Figure 5.25 Shape Constancy. The various retinal images produced by an opening door are quite different, yet we perceive a door of constant rectangular shape.

PERCEPTUAL CONSTANCIES For more Cengage Learning textbooks, visit www.cengagebrain.co.uk Courtesy Edward H. Adelson

182 CHAPTER 5 PERCEPTION

Figure 5.26 Retinal Image Size. This figure illustrates the geometric relationship between the physical size of an object and the size of its image on the retina. Arrows A and B represent objects of the same size, but one is twice as far from the eye as the other. As a result, the retinal image of A is about half the size of the retinal image of B. The object represented by arrow C is smaller than that of A, but its location closer to the eye causes it to produce a retinal image the same size as A.

Figure 5.27 Emmert's Experiment. Hold the book at normal reading distance under good light. Fixate on the cross in the center of the figure for about a minute, and then look at a distant wall. You will see an after-image of the two circles that appears larger than the stimulus. Then look at a piece of paper held close to your eyes. The afterimage will appear smaller than the stimulus. If the afterimage fades, blinking can sometimes restore it.

judgments depend on distance. Emmert used an ingenious method that involved judging the size of afterimages. Observers were first asked to fixate on the center of an image for about a minute (see Figure 5.27 for an example of such an image). Then they looked at a white screen and saw an afterimage of what they had just seen. Their task was to judge the size of the afterimage; the independent variable was how far away the screen was. Because the retinal size of the afterimage was the same regardless of the distance of the screen, any variations in judgments of the size of the afterimage had to be due to its perceived distance. When the screen was far away, the

afterimage looked large; when the screen was near, the afterimage looked small. Emmert's experiment is so easy to do that you can perform it on yourself. On the basis of such experiments, Emmert proposed that the perceived size of an object increases with both the retinal size of the object and the perceived distance of the object. This is known as the size-distance invariance principle. It explains size constancy as follows: When the distance to an object increases, the object's retinal size decreases; but if distance cues are present, perceived distance will increase. Hence, the perceived size will remain approximately constant. To illustrate: When a person walks away from you, the size of her image on your retina becomes smaller but her perceived distance becomes larger; these two changes cancel each other out, and your perception of her size remains relatively constant.

Illusions Walk into the Haunted House at Disneyland. As you nervously make your way down the first corridor, you see mask-like faces staring at you from the walls. As you move past them, the masks appear to physically swivel, ever gazing at you. Although disconcerted, you marvel at this effect, figuring that the masks must somehow be mounted on little motors that are sensitive to your approach and movement. However, in reality the masks are stationary; it is only in your perception that they move. If you somehow managed to turn on the lights and inspect the masks closely, an oddity would immediately become apparent: you are actually looking at the inside of the mask rather than the outside, as is normal. But, under the poor viewing conditions of the haunted house, you don't realize this. Your visual system makes the assumption that you are looking at a face from the outside, just as you usually do; but if this is so, it turns out that the geometry of the situation requires that you must perceive the face to be rotating as you shift position relative to it. (This is an easy demonstration that you can do for yourself. Go to a costume store and find a cheap mask - one that just goes on the front of your face, not the pull-it-down-over-your-head type. Have a friend hold the mask up across the room so that the inside of the mask is facing you. Particularly if you cover one eye, you will perceive the face as coming out at you rather than going in from you as is actually the case. Once you have that perception, you will find that as you shift back and forth, the mask will appear to rotate). The perceived-to-be-rotating mask is an example of an illusion: Your perception of something differs systematically from physical reality. The mask illusion, like many illusions, arises because of the visual system's attempts to maintain constancy - in this case its assumption that a face is, like most faces, being viewed from the outside rather than from the inside.

Constancies and illusions We have noted that the various constancies serve an important purpose: They allow us to perceive fundamental

characteristics of the world around us even when the information arriving at our sense organs (our retinas in the examples we've discussed) change dramatically as a result of different source wavelengths, different source intensities, different distances from the object or different viewing angles. For better or for worse, however, these constancies also lead to numerous optical illusions, as in the mask illusion that we have just described. The moon illusion The size-distance principle is fundamental to understanding a number of size illusions. An example is the moon illusion: When the moon is near the horizon, it looks as much as 50 percent larger than when it is high in the sky, even though in fact, the moon's retinal image is a tiny bit larger when it is directly overhead, because it is a little bit closer when directly overhead than when on the horizon (just as, for example, an airplane is closer when it is directly overhead than when you first see it on the horizon). One explanation for the moon illusion is this (see Reed, 1984; Loftus, 1985). Think about a normal flying object like an airplane that approaches you from the horizon. As we just mentioned, the geometry of the situation is that the airplane's retinal image gets larger as it moves from the

horizon to the zenith. Because an airplane is relatively close to the earth, the degree to which the retinal image gets larger is quite dramatic. Size constancy, however, compensates for this change in retinal image size in the usual fashion such that the airplane appears to remain the same physical size throughout its ascendance. Qualitatively, there is no difference between an airplane and the moon. The moon's retinal image size also (surprisingly!) increases as the moon ascends from horizon to zenith. The difference between the moon and the airplane is quantitative: the moon, unlike close-to-earth objects like airplanes that we are used to, is so far away that the change in its visual image is miniscule. However, our visual system still insists on constancy: as the moon approaches zenith, the visual system 'believes' that its retinal image size should be increasing quite a lot, just as an airplane's does. The moon's failure to increase its retinal image size in this expected manner is 'explained' by the visual system perceiving the moon's physical size to decrease; hence the moon illusion. Another way of looking at the moon illusion is that the perceived distance to the horizon is judged to be greater than the distance to the zenith. However, because the visual angle remains almost constant as the moon rises from horizon to zenith, the visual system must conclude that the moon itself is larger at the distant horizon compared to the nearer zenith (Kaufman & Rock, 1989). One way to reduce the effectiveness of the depth cues that indicate that the horizon moon is far away is to view the moon upside down. This can be done by placing your ^a HIROYUKI MATSUMOTO/GETTY IMAGES/STONE The moon looks much larger when it is near the horizon than when it is high in the sky, even though in both locations its retinal image is the same size. PERCEPTUAL CONSTANCIES For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

184 CHAPTER 5 PERCEPTION Figure 5.28 The Ames Room. A view of how the Ames room looks to an observer viewing it through the peephole. The sizes of the boy and the girl depend on which one is in the left-hand corner of the room and which one is in the right-hand corner. The room is designed to wreak havoc with our perceptions. Because of the perceived shape of the room, the relative sizes of the boy and the girl seem impossibly different. back to the moon, bending over, and viewing it through your legs. If you have a photo of the moon on the horizon, it can be done by simply turning the picture upside down (Coren, 1992). The Ames room illusion Another size illusion is created by the Ames room (named after its inventor, Adelbert Ames). Figure 5.28 shows how the Ames room looks to an observer seeing it through a peephole. When the boy is in the left-hand corner of the room (see the photograph on the left), he appears much smaller than when he is in the right-hand corner (see the photograph on the right). Yet it is the same boy in both pictures! Here we have a case in which size constancy has broken down. Why? The reason lies in the construction of the room. Although the room looks like a normal rectangular room to an observer seeing it through the peephole, it is actually shaped so that its left corner is almost twice as far away as its right corner (see the diagram in Figure 5.29). Hence, the boy on the left is much further away than the one on the right, and consequently projects a smaller retinal image. We do not correct for this difference in distance, though, because the lines in the room lead us to believe that we are looking at a normal room and therefore assume that both boys are the same distance from us. Again the visual system's only interpretation of the boy subtending a smaller angle, but being no further away is that the boy is smaller. In essence, our assumption that the room is normal blocks our application of the size-distance For more Cengage Learning textbooks, visit www.cengagebrain.co.uk invariance principle, and consequently size constancy breaks down. The 'Ames-room effect' shown in Figures 5.28 and 5.29 was used to great advantage by the movie director, Peter Jackson, in his Lord of the Rings trilogy. These movies involved different classes of

beings (e.g., Hobbits, Dwarves, Elves, and Humans) who, in keeping with J. R. R. Tolkien's original books, needed to appear to be very different sizes (e.g., Hobbits are only about half as tall as humans) even though the different beings were played by actors of similar heights. In part these effects were achieved by computer-graphics techniques, but for the most part, they were achieved by illusion. For example, Aragorn a human, would be filmed apparently walking alongside Frodo, a Hobbit. However, during the filming Viggo Mortensen playing Aragorn would be in the foreground, close to the camera, while Elijah Wood playing Frodo would actually be in the background, approximately twice as far from the camera as Mortensen. Constancies in all sensory modalities Although all the examples of constancy that we have described are visual, constancies also occur in the other senses. For example, a person will hear the same tune even if the frequencies of all its notes are doubled. Whatever the sensory modality, constancies depend on relations between features of the stimulus - between retinal size and distance in the case of size constancy, between the intensity of two adjacent regions in the case of lightness constancy, and so forth. Peephole Figure 5.29 The True Shape of the Ames Room. This figure shows the true shape of the Ames room. The boy on the left is actually almost twice as far away as the boy on the right. However, this difference in distance is not detected when the room is viewed through the peephole. (After Goldstein, 1984)

INTERIM SUMMARY | Another major function of the perceptual system is to achieve perceptual constancy - to keep the appearance of objects the same in spite of large variations in the initial representations of the stimuli received by the sense organs that are engendered by various environmental factors. | Color and brightness constancy entail perceiving the actual color and brightness of a stimulus even when the actual information arriving at the eye varies in color makeup (because of the color makeup of the ambient lighting) and in brightness (because of the level of ambient illumination) | Size constancy entails perceiving the actual size of a stimulus even when the actual size of the object's image on the retina varies because of the object's distance. | Intrinsically, constancies entail 'illusion' in the sense that by a constancy's very nature, perception differs systematically from the physical nature of the stimulus. It logically follows, and is empirically true that many visual illusions may be explained by the various constancies. | Constancies occur in all sensory modalities. | Various kinds of perceptual illusions can be explained by the perceptual system's insistence on maintaining constancies. | Although visual constancies are the most salient, constancies exist in all sensory modalities. CRITICAL THINKING QUESTIONS 1 Do you think that the moon illusion would be more pronounced if the moon were seen rising over a flat, featureless plane or if it were seen rising behind a city skyline? Suppose that you were on a boat approaching the city. Would the moon illusion be more pronounced if you were closer to the city or further from the city? 2 In what way is the behavior of a visual artist influenced by color and shape constancy? Can you think of ways in which perceptual constancies actually make the artist's task more difficult than it would be without constancy? DIVISIONS OF LABOR IN THE BRAIN In the past decade a great deal has been learned about the neural processes underlying perception. We have already touched upon some of this knowledge. In this section, we For more Cengage Learning textbooks, visit www.cengagebrain.co.uk DIVISIONS OF LABOR IN THE BRAIN will describe a bit more of what has been discovered. We will begin by talking about the neural basis of attention, and then we will turn to the visual cortex - which is a crucial waystation for incoming visual information. The neural basis of attention Recent years have produced major breakthroughs in our understanding of the neural basis of attention, particularly visual attention. The research of interest has concerned two major questions: (1) What brain structures mediate the psychological act of selecting an object to attend

to? and (2) How does the subsequent neural processing differ for attended and nonattended stimuli? Let's consider each of these questions in turn. Three brain systems in attention As previously described, there is evidence for three separate but interacting attentional systems. One functions to keep us alert. Numerous brain imaging studies have shown that when people are given tasks that require them to maintain attention on a task there is increased activity in the parietal and frontal regions of the right hemisphere of the brain. These areas are associated with the neurotransmitter norepinephrine, which is associated with arousal (Coull, Frith, Frackowiak, & Grasby, 1996). Two additional brain systems seem to mediate selective attention. The first is responsible for orienting attention to a stimulus. This system represents the perceptual features of an object, such as its location in space, its shape, and its color, and is responsible for selecting one object among many on the basis of the features associated with that object. This is sometimes referred to as the posterior system because the brain structures involved – the parietal and temporal cortex, along with some subcortical structures – are mostly located in the back of the brain (though recent research indicates a role of frontal cortex in attentional orienting). The second system, designed to control when and how these features will be used for selection, is sometimes referred to as the anterior system because the structures involved – the frontal cortex and a subcortical structure – are located in the front of the brain. In short, we can select an object for attention by focusing on its location, its shape, or its color. Although the actual selection of these features will occur in the posterior part of the brain, the selection process will be guided by the anterior part of brain. Because of this function, some researchers refer to the anterior system as the 'chief executive officer' or CEO of selective attention. Some critical findings regarding the posterior system come from PET scans of humans while they are engaged in selective-attention tasks. When observers are instructed to shift their attention from one location to another, the cortical areas that show the greatest increase in blood flow – and, hence, neural activity – are the parietal lobes of both hemispheres (Corbetta, Miezin, Shulman, &

186 CHAPTER 5 PERCEPTION Petersen, 1993). Moreover, when people with brain damage in these regions are tested on attentional tasks, they have great difficulty shifting attention from one location to another (Posner, 1988). Hence, the regions that are active when a normal brain accomplishes the task turn out to be the same areas that are damaged when a patient cannot do the task. Moreover, when single-cell recording studies are done with nonhuman primates, cells in the same brain regions are found to be active when attention must be switched from one location to the next (Wurtz, Goldberg, & Robinson, 1980). Taken together, these findings strongly indicate that activity in parietal regions of the brain mediates attending to locations. There is comparable evidence for the involvement of temporal regions in attending to the color and shape of objects (Moran & Desimone, 1985). Neural processing on attended objects Once an object has been selected for attention, what changes in neural processing occur? Consider an experiment in which a set of colored geometric objects is presented and the observer is instructed to attend only to the red ones and to indicate when a triangle is presented. The anterior system will direct the posterior system to focus on color, but what else changes in the neural processing of each stimulus? The answer is that the regions of the visual cortex that process color become more active than they would be if the observer were not selectively attending to color. More generally, the regions of the brain that are relevant to the attribute being attended to (be it color, shape, texture, motion, and so forth) will show amplified activity (Posner & Dehaene, 1994). There is also some evidence that brain regions that are relevant to unattended attributes will be inhibited (La Berge, 1995; Posner & Raichle, 1994). Some of the best evidence for this amplification of attributes that are attended to

again comes from PET studies. In one experiment (Corbetta et al., 1991), observers whose brains were being scanned viewed moving objects of varying color and form. In one condition, the individuals were instructed to detect changes among the objects in motion, while in other conditions they were instructed to detect changes among the objects in color or shape; hence, motion is the attribute attended to in the first condition, color or shape in the other conditions. As shown in Figure 5.30, even though the physical stimuli were identical in all the conditions, posterior cortical areas known to be involved in the processing of motion were found to be more active in the first condition, whereas areas involved in color or shape processing were more active in the other conditions. Attention, then, amplifies what is relevant, not only psychologically but biologically as well. For more Cengage Learning textbooks, visit www.cengagebrain.co.uk Figure 5.30 PET Images Reveal Differences in Cortical Activity. The image on the top right is from the condition in which participants attended to changes in color, whereas the images in the bottom row are from the conditions in which individuals attended to changes in shape or speed. (M. Corbetta, F. M. Miezen, S. Dobmeyer, D. L. Shulman, S. E. Persen, "Attentional Modulation of Neural Processing of Shape, Color and Velocity in Humans," *Science* V. 248 p. 1558, 1990; Reprinted by permission of the American Association for the Advancement of Science.)

The visual cortex At a general level, the part of the brain that is concerned with vision – the visual cortex – operates according to the principle of division of labor: Different regions of the visual cortex are specialized to carry out different perceptual functions (Kosslyn & Koenig, 1992; Zeki, 1993). There are over 100 million neurons in the cortex that are sensitive to visual input. Everything we know about them and the way they function has been learned through a small number of techniques. In studies involving animals, what we know is based largely on research in which electrical impulses are recorded (using microelectrodes) from single cells, as discussed in Chapter 4. Modern techniques for conducting such research owe much to the pioneering work of Hubel and Wiesel, mentioned earlier. In studies involving humans, much of what we know comes from 'natural experiments' – that is, cases of brain injury and disease that cast light on how visual behaviors relate to specific regions of the brain. Researchers in this area include neurologists (medical doctors who specialize in the brain) and neuropsychologists (psychologists who specialize in treating and studying patients with brain injury). An excellent introduction to this area is presented in Oliver Sacks's (1985) *The Man Who Mistook His Wife for a Hat*. Today the most exciting discoveries about the human brain are being made by taking pictures of the brain

Localization area Primary visual area Recognition area Figure 5.31 Two Cortical Visual Systems. The arrows going from the back of the brain toward the top depict the localization system; the arrows going from the back toward the bottom of the brain depict the recognition system. (After Mortimer Mishkin, Leslie G. Ungerleider, & Kathleen A. Macko (1983), 'Object Vision and Spatial Vision: Two Cortical Pathways,' *Trends in Neuroscience*, 6 (10):414–417.)

without surgery. This field is called brain imaging and includes techniques such as event related potentials (ERPs), positron emission tomography (PET), and functional magnetic resonance imaging (fMRI). The most important region of the brain for visual processing is the area known as the primary visual cortex, or V1. Its location at the back, or posterior, part of the brain is shown in Figure 5.31. This is the first location in the cerebral cortex to which neurons sending signals from the eye are connected. All the other visually sensitive regions of the cortex (more than 30 such locations have been identified) are connected to the eyes through V1. As has so often been the case, the function of V1 was discovered long before the development of modern recording or imaging techniques. It first became obvious when physicians examined patients who had suffered localized head injuries

through accident or war. As shown in Figure 5.32, tissue damage (technically referred to as a lesion) to a specific part of V1 was linked to blindness in very specific parts of the visual field (technically, a scotoma). Note that this form of blindness is not caused by damage to the eyes or the optic nerve; it is entirely cortical in origin. For example, the very center of the visual field – the fovea – will suffer a scotoma if a lesion occurs at the extreme rear of V1. Scotomas in more peripheral portions of the visual field are caused by lesions farther forward in V1. It is as though a map of the visual field has been stretched over the back of the cortex, with its center directly over the rearmost part of the cortex. Neurons in the primary visual cortex are sensitive to many features contained in a visual image, such as brightness, color, orientation, and motion. However, one of the most important features of these neurons is that they are each responsible for analyzing only a very tiny region of the image. In the foveal part of the image, this can be as small as less than 1 millimeter seen at arm's length. These neurons also communicate with one another only in very small regions. The benefit of this arrangement is that the entire visual field can be analyzed simultaneously and in great detail. What is missing from this analysis, however, is the ability to coordinate information that is not close together in the image – that is, to see the 'forest' in addition to the 'trees'. To accomplish this task, cortical neurons send information from V1 to the many other regions of the brain that analyze visual information. Each of these regions specializes in a particular task, such as analyzing color, motion, shape, and location. These more specialized regions are also in constant contact with V1, so that the neural communication between regions is better thought of as a conversation than as a command (Damasio, 1990; Zeki, 1993). One of the most important divisions of labor in visual analysis by the brain is between localization and recognition, to which we now turn.

Recognition versus localization systems

The idea that localization and recognition are qualitatively different tasks is supported by research findings showing that they are carried out by different regions of the visual cortex. Recognition of objects depends on a branch of the visual system that includes the primary visual cortex and a region near the bottom of the cerebral cortex. In contrast, as shown in Figure 5.31, localization of objects depends on a branch of the visual system that includes the primary visual cortex and a region of the cortex near the top of the brain. Studies with nonhuman primates show that if the recognition branch of an animal's visual system is impaired, the animal can still perform tasks that require it to perceive spatial relations between objects (one in front of the other, for example) but cannot perform tasks that require discriminating between the actual objects – for example, tasks that require discriminating a cube from a cylinder. If the location branch is impaired, the animal can perform tasks that require it to distinguish a cube from a cylinder, but it cannot perform tasks that require it to know where the objects are in relation to each other (Mishkin, Ungerleider, & Macko, 1983). Similar results have been reported in humans who have suffered parietal lobe damage, e.g., Phan, Schendel, Recanzone, and Robertson (2000). More recent research has used brain imaging to document the existence of separate object and location systems in the human brain. One widely used technique is PET (discussed in Chapter 2). A observer first has a radioactive tracer injected into her bloodstream and then is placed in a PET scanner while she performs various tasks. The scanner measures increases in radioactivity in

188 CHAPTER 5 PERCEPTION DAMAGE VISUAL FIELD LOSS a) Half-field lesion b) Local lesion c) Quarter-field lesion Occipital pole Calcarine fissure Figure 5.32 The Visual Consequences of Various Kinds of Lesions in the Primary Visual Cortex (V1). The Visual Consequences of Various Kinds of Lesions in the Primary Visual Cortex (V1) The 'map' of the visual field is upside down and mirror

reversed. Various brain regions, which indicate increases in blood flow to those regions. The regions that show the most increase in blood flow are the ones that mediate performance of the task. In one such study, observers performed two tasks, one a test of face recognition, which depends on the brain region for object recognition, and the other a test of mental rotation, which requires localization. In the face recognition task, observers saw a target picture with two test faces beneath it during each trial. One of the test faces was the face of the person depicted by the target, except for changes in orientation and lighting; the other was the face of a different person. As shown in Figure 5.33a, the observer's task was to decide which test face was the same as the target. While the observer was engaging in this task, there was an increase in blood flow in the recognition branch of the cortex (the branch terminating near the bottom of the cortex), but not in the localization branch (the branch terminating near the top of the cortex). Very different results were obtained with the mental rotation task. In this task, on each trial, observers saw a target display of a dot at some distance from a double line; beneath the target were two test displays. As shown in Figure 5.33b, one test display was the same as the target, except that it had been rotated; the other test display contained a different configuration of the dot and lines. While engaging in this task, observers showed an increase in blood flow in the localization branch of the cortex, but not in the recognition branch. Localization and recognition, therefore, are carried out in entirely different regions of the visual cortex (Grady et al., 1992; Haxby et al., 1990). The division of labor in the visual cortex does not end with the split between localization and recognition. Rather, the different kinds of information that are used in localization – eye movements, motion analysis, and depth perception, for example – are themselves processed by different subregions of the localization branch of the cortex. Similarly, the various kinds of information used in recognition – shape, color, and texture – also have specialized subregions devoted to their analysis (Livingstone & Hubel, 1988; Zeki, 1993). The upshot of all this is that the visual cortex consists of numerous 'processing modules', each of which is specialized for a particular task. The more we learn about the neural basis of other sensory modalities (and other psychological functions as well), the more this modular, or division-of-labor, approach seems to hold.

INTERIM SUMMARY | Three separate brain systems seem to mediate the psychological act of selecting an object to attend to. The first system is generally associated with arousal. The second, or posterior system, selects objects on the basis of location, shape, or color. The third, or anterior system is responsible for guiding this process, depending on the goals of the viewer. | The visual cortex operates according to the principle of division of labor. Localization is mediated by a region near the top of the cortex, and recognition by a region near the bottom of the cortex. Recognition processes are further subdivided into separate modules such as color, shape, and texture. | Recognition and localization are carried out by two different regions of the visual cortex.

Figure 5.33 Recognition and Localization Tasks. Sample items from the face-matching (left) and dot-location (right) matching tasks. (Reprinted from *Journal of Cognitive Neuroscience*, pp. 23–24, Fig. 5-2, p. 30, vol. 4:1, Winter 1992, by permission of the MIT Press, Cambridge, MA)

CRITICAL THINKING QUESTIONS 1 Why do you think the brain seems to solve many problems by dividing the work among specialized regions? What advantages may be gained by this approach? What problems might be caused by this division of labor? 2 Some people are skeptical about the value of studying perception and behavior from a biological perspective. Given what you have learned about vision and visually guided behavior, how would you argue against such skeptics?

PERCEPTUAL DEVELOPMENT An age-old question about perception is whether our abilities to

perceive are learned or innate – the familiar nature-versus-nurture problem. Contemporary psychologists no longer believe that this is an ‘either-or’ question. No one doubts that both genetics and learning influence perception; rather, the goal is to pinpoint the contribution of each and to spell out their interactions. For the modern researcher, the question ‘Must we learn to perceive?’ has given way to more specific questions: (a) What discriminatory capacities do infants have (which tells us something about inborn capacities), and For more Cengage Learning textbooks, visit www.cengagebrain.co.uk PERCEPTUAL DEVELOPMENT how does this capacity change with age under normal rearing conditions? (b) If animals are reared under conditions that restrict what they can learn (referred to as controlled stimulation), what effects does this have on their later discriminatory capacity? (c) What effects does rearing under controlled conditions have on perceptual-motor coordination? We will address each of these issues in turn. Discrimination by infants Perhaps the most direct way to find out what human perceptual capacities are inborn is to see what capacities an infant has. At first, you might think that this research should consider only newborns, because if a capacity is inborn it should be present from the first day of life. This idea turns out to be too simple, though. Some inborn capacities, such as perception of form, can appear only after other more basic capacities, such as the ability to register details, have developed. Other inborn capacities may require that there be some kind of environmental input for a certain length of time in order for the capacity to mature. Thus, the study of inborn capacities traces perceptual development from the first minute of life through the early years of childhood. Methods of studying infants It is hard for us to know what an infant perceives because it cannot talk or follow instructions, and has a fairly limited set of behaviors. To study infant perception, a researcher needs to find a form of behavior through which an infant indicates what it can discriminate. As shown in Figure 5.34, one such behavior is an infant’s tendency to look at some objects more than at others; psychologists make use of this behavior in a technique known as the preferential looking method (Teller, 1979). Two stimuli are presented to the infant side by side. The experimenter, who is hidden from the infant’s view, looks through a partition behind the stimuli and, by watching the infant’s eyes, measures the amount of time that the infant looks at each stimulus. (Usually the experimenter uses a television camera to record the infant’s viewing pattern.) During the experiment the positions of the stimuli are switched randomly. If an infant consistently looks at one stimulus more than at the other, the experimenter concludes that the infant can tell them apart – that is, discriminate between them.

A related technique is called the habituation method (Frantz, 1966; Horowitz, 1974). It takes advantage of the fact that although infants look directly at novel objects, they soon become bored with the same object – that is, they habituate. Suppose that an object is presented for a while and then replaced by a new object. To the extent that the second object is perceived as identical or highly similar to the first one, the infant should spend little time looking at it; conversely, to the extent that the second object is perceived as substantially different from the first one, the infant should spend a lot of time staring at it. By these means, an experimenter can determine whether two physical displays look the same to an infant. Using these techniques, psychologists have studied a variety of perceptual capacities in infants. Some of these capacities are needed to perceive forms, and hence are used in the task of recognition; others, particularly depth perception, are involved in the task of localization; and still others are involved in the task of keeping the appearance of perceived objects constant. Perceiving forms To be able to perceive an object, a person must first be able to discriminate one part of it from another, an ability referred to as visual acuity. Acuity is often assessed by varying both the contrast in a pattern (the difference in

brightness between dark and light regions) and the spatial frequency of the pattern (the number of times a pattern is repeated within a given area). For any particular level of contrast there are always some spatial frequencies that cannot be resolved by the visual system because they are too fine. At the other extreme, there are other spatial frequencies that cannot be seen because they change over too large an area. The method typically used in studying acuity in infants is preferential looking, with a pattern of stripes as one stimulus and a uniform gray field as the other. Initially the stripes are relatively wide, and the infant prefers to look at the pattern rather than at the uniform field. Then the researcher decreases the width of the stripes until the infant no longer shows a preference. Presumably at this point the infant can no longer discriminate a stripe from its surroundings, so that the pattern of stripes no longer has perceptible parts and looks like a uniform field. When first studied at about one month of age, infants can see some patterns, but their acuity is very low. Acuity increases rapidly over the first six months of life; then it increases more slowly, reaching adult levels between 1 and 2 years of age (Courage & Adams, 1990; Teller & Movshon, 1986). What do studies like this tell us about the infant's perceptual world? At one month, infants can distinguish among relatively large objects but cannot distinguish fine details. Such vision is sufficient to perceive some general characteristics of an object, including some of the features of a face (which create something like a pattern of dark and light stripes). Figure 5.35 uses the results of acuity experiments to simulate what one-, two-, and threemonth-old infants see when viewing a woman's face from a distance of six inches. At one month, acuity is so poor Figure 5.35 Visual Acuity and Contrast Sensitivity. Simulations of what 1-, 2-, and 3-month-old infants see when they look at a woman's face from a distance of about 6 inches and a photograph of what an adult sees. The simulations of infant perception were obtained by first determining an infant's contrast sensitivity and then applying this contrast-sensitivity function to the photograph. (From *Sensations and Perception* by E. Bruce Goldstein, © 1989, 1984, 1980 Wadsworth Publishing Co.) Figure 5.34 Testing the Visual Preferences of an infant. © LAWRENCE MIGDALE/PHOTO RESEARCHERS

CHAPTER 5 PERCEPTION For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

that it is difficult to perceive facial expressions (and indeed newborns look mostly at the outside contours of a face). By three months, acuity has improved to the point where an infant can decipher facial expressions. No wonder that infants seem so much more socially responsive at three months than at one month. Being able to discriminate dark from light edges is essential for seeing forms, but what about other aspects of object recognition? Our sensitivity to some of the shape features of objects is manifested very early in life. When presented with a triangle, even a three-day-old infant will direct its eye movements toward the edges and vertices rather than looking randomly over the form (Salapatek, 1975). Also, infants find some shapes more interesting than others. As noted in Chapter 3, they tend to look more at forms that resemble human faces, a tendency that appears to be based on a preference to attend to objects with more visual complexity in the upper portion of the object (Macchi Cassia, Turati, & Simion, 2004). By three months an infant can recognize something about the mother's face, even in a photograph, as revealed by an infant's preference to look at a photograph of the mother rather than one of an unfamiliar woman (Barrera & Maurer, 1981a). Perceiving depth Depth perception begins to appear at about three months but is not fully established until about six months. Thus, at around four months infants will begin to reach for the nearer of two objects, where nearness is signaled by binocular disparity (Granrud, 1986). A month or two later they will begin to reach for objects that are apparently nearer on the basis of monocular depth cues such as relative size, linear perspective, and shading cues (Coren, Ward, & Enns, 1999). Further evidence of the development

of monocular depth perception comes from studies using what is called a 'visual cliff', illustrated in Figure 5.36. This consists of a board placed across a sheet of glass, with a surface of patterned material located directly under the glass on the shallow side and at a distance of a few feet below the glass on the deep side. (The appearance of depth in Figure 5.36 – the 'cliff' – is created by an abrupt change in the texture gradient.) An infant who is old enough to crawl (6–7 months) is placed on the board; a patch is placed over one eye to eliminate binocular depth cues. For more Cengage Learning textbooks, visit www.cengagebrain.co.uk PERCEPTUAL DEVELOPMENT © MARTIN ROGERS/STOCK BOSTON Figure 5.36 The Visual Cliff. The 'visual cliff' is an apparatus used to show that infants and young animals are able to see depth by the time they are able to move about. The apparatus consists of two surfaces, both displaying the same checkerboard pattern and covered by a sheet of thick glass. One surface is directly under the glass; the other is several feet below it. When placed on the center board between the deep side and the shallow side, the kitten refuses to cross to the deep side but will readily move off the board onto the shallow side. (After Gibson & Walk, 1960) When the mother calls or beckons from the shallow side, the infant will consistently crawl toward her; but when the mother beckons from the deep side, the infant will not cross the 'cliff'. Thus, when an infant is old enough to crawl, depth perception is relatively well developed. Perceiving constancies Like the perception of form and depth, the perceptual constancies start to develop in the first few months of life. This is particularly true of shape and size constancy (Kellman, 1984). Consider an experiment on size constancy that used the habituation method. Four-month-old infants were first shown one teddy bear for a while and then shown a second one. The second bear was either (a) identical in physical size to the original one, but presented at a different distance so that it produced a different-sized retinal image, or (b) different in physical size from the original bear. If the infants had developed size constancy, they should perceive bear 'a' (same physical size) as identical to the one they saw originally, and hence spend little time looking at it compared to the amount of time spent looking at bear 'b' (which was actually bigger than the original). And this is exactly what happened (Granrud, 1986). Controlled stimulation We turn now to the question of how specific experiences affect perceptual capacities. To answer this question,

192 CHAPTER 5 PERCEPTION researchers have systematically varied the kind of perceptual experiences a young organism has, and then looked at the effects of this experience on subsequent perceptual performance. Absence of stimulation The earliest experiments on controlled stimulation sought to determine the effects of rearing an animal in the total absence of visual stimulation. The experimenters kept animals in the dark for several months after birth, until they were mature enough for visual testing. The idea behind these experiments was that if animals have to learn to perceive, they would be unable to perceive when first exposed to the light. The results turned out as expected: Chimpanzees that were reared in darkness for their first 16 months could detect light but could not discriminate among patterns (Riesen, 1947). However, subsequent studies showed that prolonged rearing in the dark does more than prevent learning; it causes deterioration of neurons in various parts of the visual system. It turns out that a certain amount of light stimulation is necessary to maintain the visual system. Without any light stimulation, nerve cells in the retina and visual cortex begin to atrophy (Binns & Salt, 1997; Movshon & Van Sluyters, 1981). Although these findings do not tell us much about the role of learning in perceptual development, they are important in themselves. In general, when an animal is deprived of visual stimulation from birth, the longer the period of deprivation, the greater the deficit. Adult cats, on the other hand, can have a patch over one eye for a long period without losing vision in that eye. These observations led to the idea that there is a critical period for the development of inborn

visual capacities. (A critical period is a stage in development during which the organism is optimally ready to acquire certain abilities.) Lack of stimulation during a critical period for vision can permanently impair the visual system (Cynader, Timney, & Mitchell, 1980). Limited stimulation Researchers no longer deprive animals of stimulation for a long time; instead, they study the effects of rearing animals that receive stimuli in both eyes, but only certain kinds of stimuli. Researchers have raised kittens in an environment in which they see only vertical stripes or only horizontal stripes. The kittens become blind to stripes in the orientation – horizontal or vertical – that they do not experience. And single-cell recording studies show that many cells in the visual cortex of a ‘horizontally reared’ cat respond to horizontal stimuli and none responds to vertical stimuli, whereas the opposite pattern is found in the visual cortex of ‘vertically reared’ cat (Blake, 1981; Movshon & Van Sluyters, 1981). This blindness seems to be caused by the degeneration of cells in the visual cortex. For more Cengage Learning textbooks, visit www.cengagebrain.co.uk Of course, researchers do not deprive humans of normal visual stimulation, but sometimes this happens naturally or as a consequence of medical treatment. For example, after eye surgery the eye that was operated on is usually covered with a patch. If this happens to a child in the first year of life, the acuity of the patched eye is reduced (Awaya et al., 1973). This suggests that there is a critical period early in the development of the human visual system similar to that in animals; if stimulation is restricted during this period, the system will not develop normally. The critical period is much longer in humans than in animals. It may last as long as eight years, but the greatest vulnerability occurs during the first two years of life (Aslin & Banks, 1978). None of these facts indicates that we have to learn to perceive. Rather, the facts show that certain kinds of stimulation are essential for the maintenance and development of perceptual capacities that are present at birth. But this does not mean that learning has no effect on perception. For evidence of such effects, we need only consider our ability to recognize common objects. The fact that we can recognize a familiar object more readily than an unfamiliar one – a dog versus an aardvark, for example – must certainly be due to learning. If we had been reared in an environment rich in aardvarks and sparse in dogs, we could have recognized the aardvark more readily than the dog.

Active perception When it comes to coordinating perceptions with motor responses, learning plays a major role. The evidence for this comes from studies in which observers receive normal stimulation but are prevented from making normal responses to that stimulation. Under such conditions, perceptual-motor coordination does not develop. For example, in one classic study, two kittens that had been reared in darkness had their first visual experience in the ‘kitten carousel’ illustrated in Figure 5.37. As the active kitten walked, it moved the passive kitten riding in the carousel. Although both kittens received roughly the same visual stimulation, only the active kitten had this stimulation produced by its own movement. And only the active kitten successfully learned sensory-motor coordination; for example, when picked up and moved toward an object, only the active kitten learned to put out its paws to ward off a collision. Similar results have been obtained with humans. In some experiments, people wear prism goggles that distort the directions of objects. Immediately after putting on these goggles, they temporarily have trouble reaching for objects and often bump into things. If they move about and attempt to perform motor tasks while wearing the goggles, they learn to coordinate their movements with the actual location of objects rather than with their apparent locations. On the other hand, if a person is pushed in a wheelchair he or she does not adapt to the

Figure 5.37 The Importance of Self-Produced Movements. Both kittens received roughly the same visual stimulation, but only the active kitten had this stimulation produced by its own movement.

(R. Held and A. Held (1963) 'Movement Produced in the Development of Visually Guided Behavior', from *Journal of Comparative and Physiological Psychology*, 56:872-876. Copyright © 1963 by the American Psychological Association. Adapted with permission.) goggles. Apparently, self-produced movement is essential to prism adaptation (Held, 1965). In sum, the evidence indicates that we are born with considerable perceptual capacities. The natural development of some of these capacities may require years of normal input from the environment. But there clearly are learning effects on perception as well; these are particularly striking when perception must be coordinated with motor behavior. This chapter, like the preceding one, includes many examples of the interplay between psychological and biological approaches. Throughout the chapter we have encountered cases in which specific psychological functions are implemented by specific cells or brain regions. We have seen that specialized cells are used to perceive motion and that separate parts of the brain are used to register the visual features of location, shape, and color. Still other regions of the brain are involved in determining which of these features will be used to control behaviors and actions. These and other examples illustrate how significant the findings of biological research can be in the study of psychological processes. For more Cengage Learning textbooks, visit

www.cengagebrain.co.uk PERCEPTUAL DEVELOPMENT Figure 5.38 After looking at this picture, look back at the left panel of Figure 5-1 (page 152). Now what do you see? INTERIM SUMMARY I

Research on perceptual development is concerned with the extent to which perceptual capacities are inborn and the extent to which they are learned through experience. I To determine inborn capacities, researchers study the discrimination capacities of infants with methods such as preferential looking and habituation. Perceptual constancies begin to develop as early as six months. I Animals raised in darkness suffer permanent visual impairment, and animals raised with a patch over one eye become blind in that eye, suggesting a critical period early in life when lack of normal stimulation produces deficiency in an innate perceptual capacity. CRITICAL THINKING QUESTION 1 Do you think that in general infants are more or less able to perceive the world than their parents think they are?

SEEING BOTH SIDES IS PERCEPTUAL DEVELOPMENT AN INNATE OR SOCIALLY ACQUIRED PROCESS?

Perceptual development is an intrinsic process Elizabeth S. Spelke, Massachusetts Institute of Technology Human beings have a striking capacity to learn from one another. This capacity already is evident in the 1-year-old child, who can learn the meaning of a new word by observing just a few occasions of its use and who can learn the functions of a new object simply by watching another person act on it. The rapid and extensive learning that occurs in early childhood suggests that much of what humans come to know and believe is shaped by our encounters with other things and people. But is our very ability to perceive things and people itself the result of learning? Or, does perception originate in intrinsically generated growth processes and develop in relative independence of one's encounters with things perceived? For two millennia, most of the thinkers who have pondered this question have favored the view that humans learn to perceive, and that the course of development proceeds from meaningless, unstructured sensations to meaningful, structured perceptions. Research on human infants nevertheless provides evidence against this view. For example, we now know that newborn infants perceive depth and use depth information as adults do, to apprehend the true sizes and shapes of objects. Newborn infants divide the speech stream into the same kinds of sound patterns as do adults, focusing in particular on the set of sound contrasts used by human languages. Newborn infants distinguish human faces from other patterns and orient to faces preferentially. Finally, newborn infants are sensitive to many of the features of objects that adults use to distinguish one thing from another, and they appear to

combine featural information in the same kinds of ways as do adults. How does perception change after the newborn period? With development, infants have been found to perceive depth, objects, and faces with increasing precision. Infants also come to focus on the speech contrasts that are relevant to their own language in preference to speech contrasts relevant to other languages. (Interestingly, this focus appears to result more from a decline in sensitivity to foreign language contrasts than from an increase in sensitivity to native language contrasts.) Finally, infants become sensitive to new sources of information about the environment, such as stereoscopic information for depth, configural information for object boundaries, and new reference frames for locating objects and events. These developments bring greater precision and richness to infants' perceptual experience, but they do not change the infant's world from a meaningless flow of sensation to a meaningful, structured environment. The findings from studies of human infants gain further support from studies of perceptual development in other animals. Since the pioneering work of Gibson and Walk, we have known that depth perception develops without visual experience in every animal tested: Innate capacities for perceiving depth allow newborn goats to avoid falling off cliffs, and they allow darkreared rats and cats to avoid bumping into approaching surfaces. More recent studies reveal that newborn chicks perceive the boundaries of objects much as human adults do, and they even represent the continued existence of objects that are hidden. Studies of animals' developing brains reveal that both genes and intrinsically structured neural activity are crucial to the development of normally functioning perceptual systems, but encounters with the objects of perception — external things and events — play a much lesser role. As with human infants, normal visual experience enriches and attunes young animals' perceptual systems, and abnormal visual experience may greatly perturb their functioning. Like human infants, however, other animals do not need visual experience to transform their perceptual world from a flow of unstructured sensations into a structured visual layout. In sum, perception shows considerable structure at birth and continuity over development. This continuity may help to explain why young human infants are so adept at learning from other people. Consider an infant who watches an adult twist a lid off a jar while saying, 'Let's open it'. If the infant could not perceive the lid and jar as distinct movable and manipulable objects, she would not be able to make sense of the adult's action. If she could not perceive the sounds that distinguish 'open' from other words, she could not begin to learn about this distinctive utterance. And if she could not perceive the person as an agent in some way like herself, then watching the person's action and listening to his speech would reveal nothing about what the infant herself could learn to do or say. Infants' prodigious abilities to learn, therefore, may depend critically on equally prodigious, unlearned abilities to perceive.

CHAPTER 5 PERCEPTION For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

SEEING BOTH SIDES IS PERCEPTUAL DEVELOPMENT AN INNATE OR SOCIALLY ACQUIRED PROCESS?

Perceptual development is an activity-dependent process: Mark Johnson, University of London Most developmental scientists now agree that both nature and nurture are essential for the normal development of perception. However, there is still much dispute about the extent to which either nature or nurture is the more important factor. Points of view on this issue are more than just philosophical musings; they affect the kinds of research programs that are undertaken. Since the 1980s a major thrust in developmental psychology has centered on identifying and delineating aspects of perceptual and cognitive function that can be termed innately specified core knowledge (Spelke & Kinzler, 2007). Core knowledge is contrasted with learning mechanisms engaged by visual experience. I argue here that this line of thinking fails to reflect the fact that the most interesting phenomena in development involve interactions between acquired and intrinsic

processes, and that common mechanisms of brain adaptation may underlie the two processes. I propose that perceptual development is better characterized as an activity-dependent process involving complex and subtle interactions at many levels, and that the infant actively seeks out the experience it needs for its own further brain development. To begin to illustrate my point, let's consider neurobiological work on the prenatal development of the visual cortex in another species, rodents. The neurons studied in these experiments are those involved in binocular vision. Experiments show that the prenatal tuning of these neurons arises through their response to internally generated waves of electrical activity from the main inputs to the visual cortex, the lateral geniculate nucleus and eye (Katz & Shatz, 1996). In other words, the response properties of these visual cortical neurons are shaped by a kind of 'virtual environment' generated by cells elsewhere in the brain and eye. Although the term innate can be stretched to cover this example of development, we could equally well describe this process as the cortical cells learning from the input provided by their cousins in the LGN and eye. Further, after birth the same cortical neurons continue to be tuned in the same way, except that now their input also reflects the structure of the world outside the infant. Thus, when we examine development in detail, it becomes harder to argue that 'innate knowledge' is fundamentally different from learning. Another example of the role of activity-dependent processes in perceptual development comes from the ability to detect and recognize faces. Because regions of the adult human cortex are specialized for processing faces, some have argued that this ability is innate. However, experiments with infants reveal a more complex story (Johnson, 2005). The tendency for newborns to look more toward faces turns out to be based on a very primitive reflex-like system that may be triggered by a stimulus as simple as three high-contrast blobs in the approximate locations of the eyes and mouth. This simple attention bias, together with a sensitivity to the human voice, is sufficient to ensure that newborns look much more at faces than at other objects and patterns over the first weeks of life. One consequence of this is that developing brain circuits on the visual recognition pathway of the cortex get more input related to faces and thus are shaped by experience with this special type of visual stimulus. We can now study this process by using new brain-imaging methods. Such studies have shown that the brains of young children show less localized and less specialized processing of faces in the cortex than do the brains of adults. It is not until around 10 years old that children start to show the same patterns of brain specialization for processing faces as adults, by which time they have had as much as 10,000 hours of experience of human faces. Another example comes from the study of infants' eye movements to visual targets. Although newborns are capable of some primitive reflexive eye movements, only much later in the first year can they make most of the kinds of complex and accurate saccades seen in adults. One view is that the very limited ability present in newborns is just sufficient to allow them to practice and develop new brain circuits for the more complex integration of visual and motor information necessary for adultlike eye movements. And practice they do! Even by four months, babies have already made more than 3 million eye movements. Once again, it appears that infants actively contribute to their own subsequent development. These considerations should also make us skeptical about claims made for innate perceptual abilities based on experiments with babies that are several months old. In fact, when the same experiments were done with younger infants, quite different results have sometimes been obtained, suggesting dramatic changes in perceptual abilities over the first few weeks and months after birth (Haith, 1998). Infants are not passively shaped by either their genes or their environment. Rather, perceptual development is an activity-dependent process in which, during postnatal life, the infant plays an active role in generating the experience it needs for its own subsequent brain development. Mark Johnson PERCEPTUAL DEVELOPMENT For more Cengage

196 CHAPTER 5 PERCEPTION CHAPTER SUMMARY The study of perception deals with the question of how organisms process and organize incoming raw, sensory information in order to (a) form a coherent representation or model of the world within which the organism dwells and (b) use that representation to solve naturally occurring problems, such as navigating, grasping, and planning. Five major functions of the perceptual system are: (a) Determining which part of the sensory environment to attend to, (b) localizing, or determining where objects are, (c) recognizing, or determining what objects are, (d) abstracting the critical information from objects, and (e) keeping appearance of objects constant, even though their retinal images are changing. Another area of study is how our perceptual capacities develop. Selective attention is the process by which we select some stimuli for further processing while ignoring others. In vision, the primary means of directing our attention are eye movements. Most eye fixations are on the more informative parts of a scene. Selective attention also occurs in audition. Usually we are able to selectively listen by using cues such as the direction from which the sound is coming and the voice characteristics of the speaker. Our ability to selectively attend is mediated by processes that occur in the early stages of recognition as well as by processes that occur only after the message's meaning has been determined. To localize objects we must first separate them from one another and then organize them into groups. These processes were first studied by Gestalt psychologists, who proposed several principles of organization. One such principle is that we organize a stimulus into regions corresponding to figure and ground. Other principles concern the bases that we use to group objects together, including proximity, closure, similarity, good continuation, and closure. Localizing an object requires that we know its distance from us. This form of perception, known as depth perception, is usually thought to be based on depth cues. Monocular depth cues include relative size, interposition, relative height, linear perspective, shading, and motion parallax. A binocular depth cue is binocular disparity, which results from the fact that any object produces slightly different images on the two retinas. Localizing an object sometimes requires that we know the direction in which an object is moving. Motion perception can be produced in the absence of an object moving across our retina. One example of this phenomenon is stroboscopic motion, in which a rapid series of still images induces apparent movement; another example is induced motion, in which movement of a large object induces apparent movement of a smaller stationary object. Perception of real motion (movement of a real object through space) is implemented by specific cells in the visual system, as indicated by single-cell recordings and experiments on selective adaptation. Recognizing an object requires that the various features associated with the object (e.g., shapes, colors) be correctly bound together. It is generally believed that attention is required for this binding process; when such binding fails, an illusory conjunction – the incorrect conjunction of two or more features of different objects – may occur. Recognizing an object amounts to assigning it to a category and is based mainly on the shape of the object. In early stages of recognition, the visual system uses retinal information to describe the object in terms of features like lines and angles; neurons that detect such features (feature detectors) have been found in the visual cortex. In later stages of recognition, the system matches the description of the object with shape descriptions stored in memory to find the best match. Matching can be explained by a connectionist model or network. The bottom level of the network contains features and the next level contains letters; an excitatory connection between a feature and a letter means that the feature is part of a letter, while an inhibitory connection means that the feature is not part of the letter. When a letter is presented, it

activates some features in the network, which pass their activation or inhibition up to letters; the letter that receives the most activation is the best match to the input. The network can be expanded to include a level of words and to

explain why a letter is easier to recognize when presented in a word than when presented alone. The shape features of natural objects are more complex than lines; they are similar to simple geometric forms such as cylinders, cones, blocks, and wedges. A limited set of such forms may be sufficient in combination to describe the shapes of all objects that people can recognize. Research indicates that face recognition involves processes separate from object recognition. Object recognition depends on processing features, and face recognition depends in part on processing overall configuration. Bottom-up recognition processes are driven solely by the input, whereas top-down recognition processes are driven by a person's knowledge and expectations. Top-down processes underlie context effects in perception: The context sets up a perceptual expectation, and when this expectation is satisfied, less input information than usual is needed for recognition. Another major function of the perceptual system is to achieve perceptual constancy – that is, to keep the appearance of objects the same in spite of large changes in the stimuli received by the sense organs. Lightness constancy refers to the fact that an object appears equally light regardless of how much light it reflects, and color constancy means that an object looks roughly the same color regardless of the light source illuminating it. In both cases, constancy depends on relations between the object and elements of the background. Two other well-known perceptual constancies are shape and location constancy. Size constancy refers to the fact that an object's apparent size remains relatively constant no matter how far away it is. The perceived size of an object increases with both the retinal size of the object and the perceived distance of the object, in accordance with the size-distance invariance principle. Thus, as an object moves away from the perceiver, the size of its retinal image decreases but the perceived distance increases, and the two changes cancel each other out, resulting in constancy. This principle can be used to explain certain kinds of perceptual illusions. For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

CHAPTER SUMMARY 15 Two separate brain systems seem to mediate the psychological act of selecting an object to attend to. In the posterior system, objects are selected on the basis of location, shape, or color. The anterior system is responsible for guiding this process, depending on the goals of the viewer. PET studies further show that once an object has been selected, activity is amplified in the posterior regions of the brain that are relevant to the attribute being attended to. The visual cortex operates according to the principle of division of labor. Localization and recognition are carried out by different regions of the brain, with localization mediated by a region near the top of the cortex and recognition by a region near the bottom of the cortex. Recognition processes are further subdivided into separate modules: for example, color, shape, and texture. Research on perceptual development is concerned with the extent to which perceptual capacities are inborn and the extent to which they are learned through experience. To determine inborn capacities, researchers study the discrimination capacities of infants using methods such as preferential looking and habituation. Acuity, which is critical to recognition, increases rapidly during the first six months of life and then increases more slowly. Depth perception begins to appear at about three months but is not fully established until about six months. Perceptual constancies begin to develop as early as six months. Animals raised in darkness suffer permanent visual impairment, and animals raised with a patch over one eye become blind in that eye. Adult animals do not lose vision even when deprived of stimulation for long periods. These results suggest that there is a critical period early in life during which lack of normal stimulation produces deficiency in

an innate perceptual capacity. If stimulation early in life is controlled in such a way that certain kinds of stimuli are absent, both animals and people become insensitive to the stimuli of which they have been deprived; again, this effect does not have much to do with learning. Perceptual-motor coordination must be learned, however. Both animals and people require self-produced movement to develop normal coordination.

CORE CONCEPTS perception symbol theory of ecological optics model of the environment perceptual constancy eye fixations saccade weapon focus shadowing inattention blindness change blindness depth cues binocular disparity stroboscopic motion selective adaptation selective attention primitive features binding problem illusory conjunction feature-integration theory visual search task dynamic control theory simple cell complex cell hypercomplex cell connectionist models node object recognition network excitatory connections augmented network top-down feedback connections geons bottom-up versus top-down processes mcgurk effect prosopagnosia impression effect agnosia associative agnosia abstraction constancy available wavelengths source wavelengths spatial localization reflectance characteristic illusion posterior system anterior system preferential looking method habituation method WEB RESOURCES <http://www.atkinsonhilgard.com/> Take a quiz, try the activities and exercises, and explore web links. <http://www.yorku.ca/eye/thejoy.htm> Click on Fun Things in Vision and tease your senses while you learn more about perception. Then explore perception-specific topics like size perception, shape constancy, and more. <http://www.exploratorium.edu/imagery/exhibits> Some more examples of illusions can be found on this site from the Exploratorium in San Francisco. http://psych.hanover.edu/Krantz/sen_tut.html This site offers you a collection of tutorials related to sensation and perception. CHAPTER 5 PERCEPTION For more Cengage Learning textbooks, visit www.cengagebrain.co.uk

CD-ROM LINKS Psyk.Trek 3.0 Check out CD Unit 3, Sensation and Perception 3d Perception of color 3e Gestalt psychology 3f Depth perception 3g Visual illusions For more Cengage Learning textbooks, visit www.cengagebrain.co.uk CD-ROM LINKS 199

Revision #1

Created 2026-01-04 19:12:14 UTC by Omar Ayman

Updated 2026-01-04 19:12:14 UTC by Omar Ayman