

SOME QUESTIONS WE WILL CONSIDER

- ▶ What is cognitive neuroscience, and why is it necessary? (26)
- ▶ How is information transmitted from one place to another in the nervous system? (29)
- ▶ How are things in the environment, such as faces and places, represented in the brain? (42)
- ▶ What are neural networks, and what is their role in cognition? (45)

As we discussed in Chapter 1, research on the mind has been on a roller-coaster ride that began with a promising start in the 1800s with the early research of Donders and Ebbinghaus, only to be derailed by Watson's behaviorism in the early 1900s and Skinner's operant conditioning in the 1930s. Finally, in the 1950s and 1960s, clearer minds decided that it was important to return to the study of the mind and began doing experiments based on the information-processing model that was inspired by digital computers.

But just as this cognitive revolution was beginning, something else was happening that would have a huge impact on our understanding of the mind. In the 1950s, a number of research papers began appearing that involved recording nerve impulses from single neurons. As we will see, research studying the relationship between neural responding and cognition began long before the 1950s, but technological advances led to a large increase in physiological research beginning just about the same time the cognitive revolution was happening.

In this chapter, we take up the story of **cognitive neuroscience**—the study of the physiological basis of cognition. We begin by discussing the idea of “levels of analysis,” which is our rationale behind studying the physiology of the mind, and we then go back in time to the 19th and early 20th century to look at the early research that set the stage for amazing discoveries that were to be made beginning in the 1950s.

▶ Levels of Analysis

Levels of analysis refers to the idea that a topic can be studied in a number of different ways, with each approach contributing its own dimension to our understanding. To understand what this means, let's consider a topic outside the realm of cognitive psychology: understanding the automobile.

Our starting point for this problem might be to take a car out for a test drive. We could determine its acceleration, its braking, how well it corners, and its gas mileage. When we have measured these things, which come under the heading of “performance,” we will know a lot about the particular car we are testing. But to learn more, we can consider another level of analysis: what is going on under the hood. This would involve looking at the mechanisms responsible for the car's performance: the motor and the braking and steering systems. For example, we can describe the car as being powered by a four-cylinder 250 HP internal combustion engine and having independent suspension and disc brakes.

But we can look even deeper into the operation of the car by considering another level of analysis designed to help us understand how the car's engine works. One approach would be to look at what happens inside a cylinder. When we do this, we see that when vaporized gas enters the cylinder and is ignited by the spark plug, an explosion occurs that pushes the cylinder down and sends power to the crankshaft and then to the wheels. Clearly, considering the automobile from the different levels of driving the car, describing the motor, and observing what happens inside a cylinder provides more information about cars than simply measuring the car's performance.

Applying this idea of levels of analysis to cognition, we can consider measuring behavior to be analogous to measuring the car's performance, and measuring the physiological processes behind the behavior as analogous to what we learned by looking under the hood. And just as we can study what is happening under a car's hood at different levels, we can

study the physiology of cognition at levels ranging from the whole brain, to structures within the brain, to chemicals that create electrical signals within these structures.

Consider, for example, a situation in which Gil is talking with Mary in the park (Figure 2.1a), and then a few days later he passes the park and remembers what she was wearing and what they talked about (Figure 2.1b). This is a simple behavioral description of having an experience and later having a memory of that experience.

But what is going on at the physiological level? During the initial experience, in which Gil perceives Mary as he is talking with her, chemical processes occur in Gil's eyes and ears, which create electrical signals in neurons (which we will describe shortly); individual brain structures are activated, then multiple brain structures are activated, all leading to Gil's perception of Mary and what is happening as they talk (Figure 2.1a).

Meanwhile, other things are happening, both during Gil's conversation with Mary and after it is over. The electrical signals generated as Gil is talking with Mary trigger chemical and electrical processes that result in the storage of Gil's experiences in his brain. Then, when Gil passes the park a few days later, another sequence of physiological events is triggered that retrieves the information that was stored earlier, which enables him to remember his conversation with Mary (Figure 2.2b).

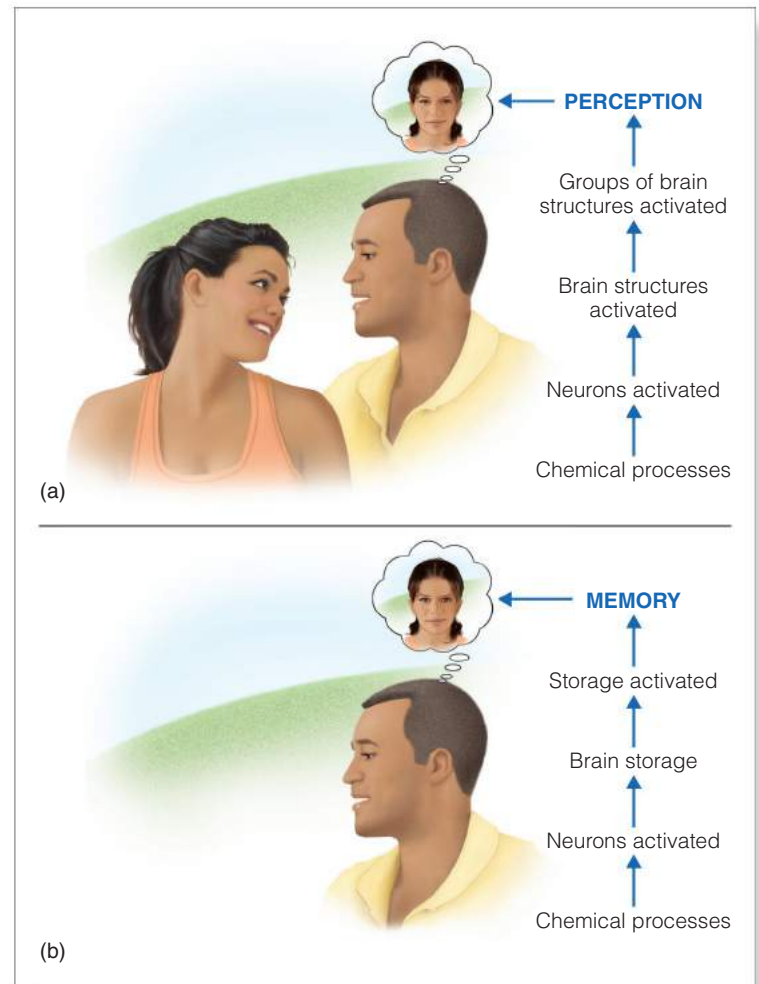
We have gone a long way to make a point, but it is an important one. To fully understand any phenomenon, whether it is how a car operates or how people remember past experiences, it needs to be studied at different levels of analysis. In this book, we will be describing research in cognition at both the behavioral and physiological levels. We begin our description of physiology by considering one of the basic building blocks of the nervous system: the neuron.

► Neurons: Basic Principles


How is it possible that the 3.5-pound structure called the brain could be the seat of the mind? The brain appears to be static tissue. Unlike the heart, it has no moving parts. Unlike the lungs, it doesn't expand or contract. And when observed with the naked eye, the brain looks almost solid. As it turns out, to understand the relation between the brain and the mind—and specifically to understand the physiological basis for everything we perceive, remember, and think—it is necessary to look within the brain and observe the small units called **neurons** that create and transmit information about what we experience and know.

Early Conceptions of Neurons

For many years, the nature of the brain's tissue was a mystery. Looking at the interior of the brain with the unaided eye gives no indication that it is made up of billions of smaller units. But in the 19th century, anatomists applied special stains to brain tissue, which increased



► **Figure 2.1** Physiological levels of analysis. (a) Gil perceives Mary and their surroundings as he talks with her. The physiological processes involved in Gil's perception can be described at levels ranging from chemical reactions to single neurons, to structures in the brain, to groups of structures in the brain. (b) Later, Gil remembers his meeting with Mary. The physiological processes involved in remembering can also be described at different levels of analysis.



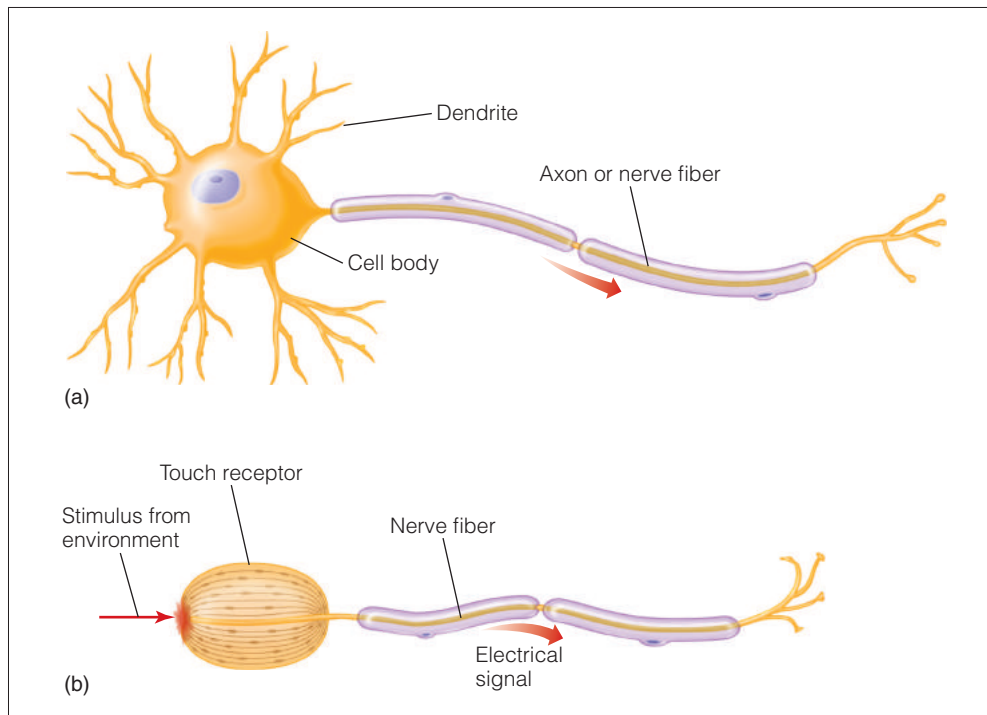
► **Figure 2.2** (a) Nerve net theory proposed that signals could be transmitted throughout the net in all directions. (b) A portion of the brain that has been treated with Golgi stain shows the shapes of a few neurons. The arrow points to a neuron's cell body. The thin lines are dendrites or axons (see **Figure 2.3**).

the contrast between different types of tissue within the brain. When they viewed this stained tissue under a microscope, they saw a network they called a **nerve net** (**Figure 2.2a**). This network was believed to be continuous, like a highway system in which one street connects directly to another, but without stop signs or traffic lights. When visualized in this way, the nerve net provided a complex pathway for conducting signals uninterrupted through the network.

One reason for describing the microstructure of the brain as a continuously interconnected network was that the staining techniques and microscopes used during that period could not resolve small details, and without these details the nerve net appeared to be continuous. However, in the 1870s, the Italian anatomist Camillo Golgi (1843–1926) developed a staining technique in which a thin slice of brain tissue was immersed in a solution of silver nitrate. This technique created pictures like the one in **Figure 2.2b**, in which fewer than 1 percent of the cells were stained, so they stood out from the rest of the tissue. (If all of the cells had been stained, it would be difficult to distinguish one cell from another because the cells are so tightly packed.) Also, the cells that were stained were stained completely, so it was possible to see their structure.

Meanwhile, the Spanish physiologist Ramon y Cajal (1852–1934) was using two techniques to investigate the nature of the nerve net. First, he used the Golgi stain, which stained only some of the cells in a slice of brain tissue. Second, he decided to study tissue from the brains of newborn animals, because the density of cells in the newborn brain is small compared with the density in the adult brain. This property of the newborn brain, combined with the fact that the Golgi stain affects less than 1 percent of the neurons, made it possible for Cajal to clearly see that the nerve net was not continuous but was instead made up of individual units connected together (Kandel, 2006). Cajal's discovery that individual units called neurons were the basic building blocks of the brain was the centerpiece of **neuron doctrine**—the idea that individual cells transmit signals in the nervous system, and that these cells are not continuous with other cells as proposed by nerve net theory.

Figure 2.3a shows the basic parts of a neuron. The **cell body** is the metabolic center of the neuron; it contains mechanisms to keep the cell alive. The function of **dendrites** that branch out from the cell body is to receive signals from other neurons. **Axons** (also called **nerve fibers**) are usually long processes that transmit signals to other neurons. **Figure 2.3b** shows a neuron with a receptor that receives stimuli from the environment—pressure, in this example. Thus, the neuron has a receiving end and a transmitting end, and its role, as visualized by Cajal, was to transmit signals.




► **Figure 2.3** (a) Basic components of a neuron in the cortex. (b) A neuron with a specialized receptor in place of the cell body. This receptor responds to pressure on the skin.

Cajal also came to some other conclusions about neurons: (1) There is a small gap between the end of a neuron's axon and the dendrites or cell body of another neuron. This gap is called a **synapse** (Figure 2.4). (2) Neurons are not connected indiscriminately to other neurons but form connections only to specific neurons. This forms groups of interconnected neurons, which together form **neural circuits**. (3) In addition to **neurons in the brain**, there are also neurons that are specialized to pick up information from the environment, such as the neurons in the eye, ear, and skin. These neurons, called **receptors** (Figure 2.3b), are similar to brain neurons in that they have an axon, but they have specialized receptors that pick up information from the environment.

Cajal's idea of individual neurons that communicate with other neurons to form neural circuits was an enormous leap forward in the understanding of how the nervous system operates. The concepts introduced by Cajal—individual neurons, synapses, and neural circuits—are basic principles that today are used to explain how the brain creates cognitions. These discoveries earned Cajal the Nobel Prize in 1906, and today he is recognized as “the person who made this cellular study of mental life possible” (Kandel, 2006, p. 61).

The Signals That Travel in Neurons

Cajal succeeded in describing the structure of individual neurons and how they are related to other neurons, and he knew that these neurons transmitted signals. However, determining the



► **Figure 2.4** (a) Neuron synapsing on the cell body of another neuron. (b) Close-up of the synapse showing the space between the end of one neuron and the cell body of the next neuron, and neurotransmitter being released.

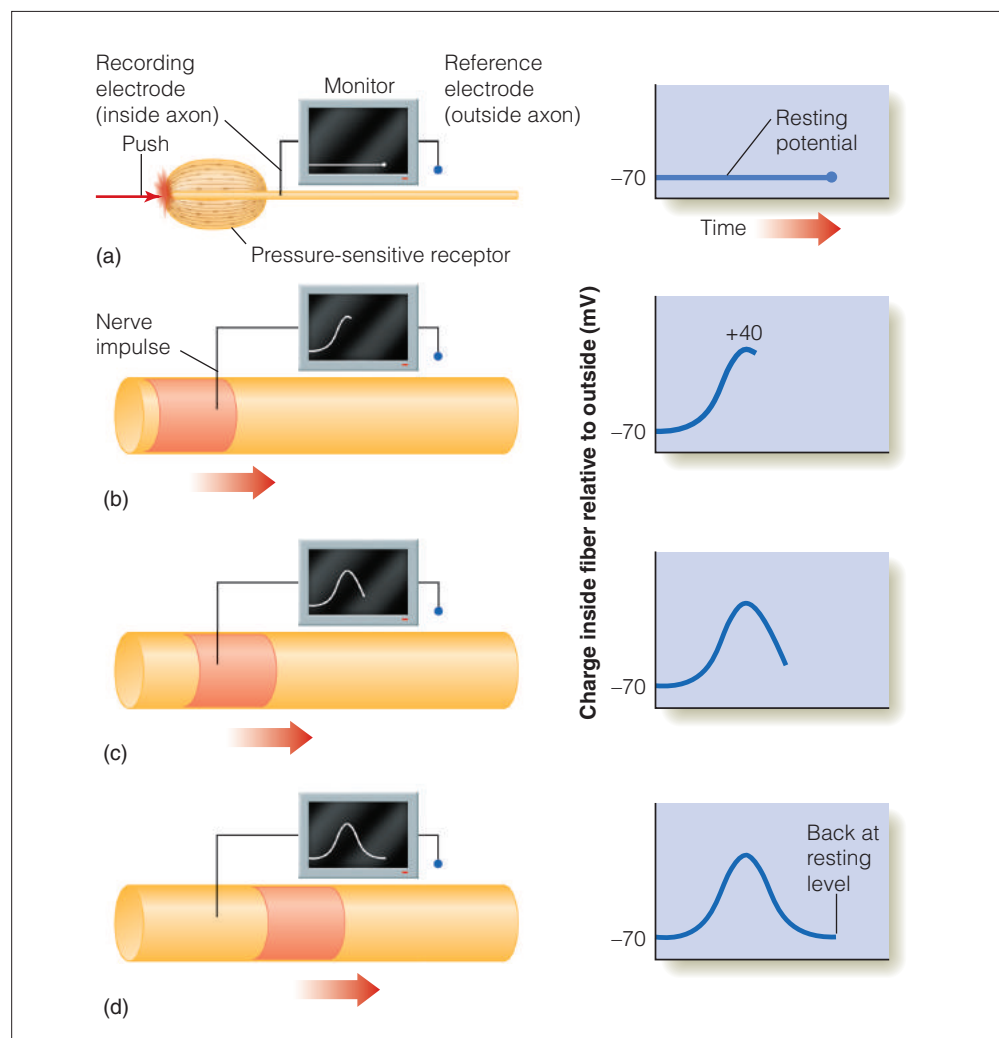
exact nature of these signals had to await the development of electronic amplifiers that were powerful enough to make the extremely small electrical signals generated by the neuron visible. In the 1920s, Edgar Adrian was able to record electrical signals from single sensory neurons, an achievement for which he was awarded the Nobel Prize in 1932 (Adrian, 1928, 1932).

METHOD Recording from a Neuron

Adrian recorded electrical signals from single neurons using **microelectrodes**—small shafts of hollow glass filled with a conductive salt solution that can pick up electrical signals at the electrode tip and conduct these signals back to a recording device. Modern physiologists use metal microelectrodes.

Figure 2.5 shows a typical setup used for recording from a single neuron. There are two electrodes: a **recording electrode**, shown with its recording tip inside the neuron,¹ and a **reference electrode**, located some distance away so it is not affected by the electrical signals. The difference in charge between the recording and reference electrodes is fed into a computer and displayed on the computer's screen.

► **Figure 2.5** Recording an action potential as it travels down an axon. (a) When the nerve is at rest, there is a difference in charge, called the *resting potential*, of -70 millivolts (mV) between the inside and outside of the axon. The difference in charge between the recording and reference electrodes is fed into a computer and displayed on a computer monitor. This difference in charge is displayed on the right. (b) As the nerve impulse, indicated by the red band, passes the electrode, the inside of the fiber near the electrode becomes more positive. (c) As the nerve impulse moves past the electrode, the charge in the fiber becomes more negative. (d) Eventually the neuron returns to its resting state.



¹In practice, most recordings are achieved with the tip of the electrode positioned just outside the neuron because it is technically difficult to insert electrodes into the neuron, especially if it is small. However, if the electrode tip is close enough to the neuron, the electrode can pick up the signals generated by the neuron.

When the axon, or nerve fiber, is at rest, the meter records a difference in potential between the tips of the two electrodes of -70 millivolts (a millivolt is $1/1000$ of a volt), as shown on the right in **Figure 2.5a**. This value, which stays the same as long as there are no signals in the neuron, is called the **resting potential**. In other words, the inside of the neuron has a charge that is 70 mV more negative than the outside, and this difference continues as long as the neuron is at rest.

Figure 2.5b shows what happens when the neuron's receptor is stimulated so that a **nerve impulse** is transmitted down the axon. As the impulse passes the recording electrode, the charge inside the axon rises to $+40$ millivolts, compared to the outside. As the impulse continues past the electrode, the charge inside the fiber reverses course and starts becoming negative again (**Figure 2.5c**), until it returns to the resting potential (**Figure 2.5d**). This impulse, which is called the **action potential**, lasts about 1 millisecond ($1/1000$ of a second).

Figure 2.6a shows action potentials on a compressed time scale. Each vertical line represents an action potential, and the series of lines indicates that a number of action potentials are traveling past the electrode. **Figure 2.6b** shows one of the action potentials on an expanded time scale, as in **Figure 2.5**. There are other electrical signals in the nervous system, but we will focus here on the action potential because it is the mechanism by which information is transmitted throughout the nervous system.


In addition to recording action potentials from single neurons, Adrian made other discoveries as well. He found that each action potential travels all the way down the axon without changing its height or shape. This property makes action potentials ideal for sending signals over a distance, because it means that once an action potential is started at one end of an axon, the signal will still be the same size when it reaches the other end.

At about the same time Adrian was recording from single neurons, other researchers were showing that when the signals reach the synapse at the end of the axon, a chemical called a **neurotransmitter** is released. This neurotransmitter makes it possible for the signal to be transmitted across the gap that separates the end of the axon from the dendrite or cell body of another neuron (see **Figure 2.4b**).

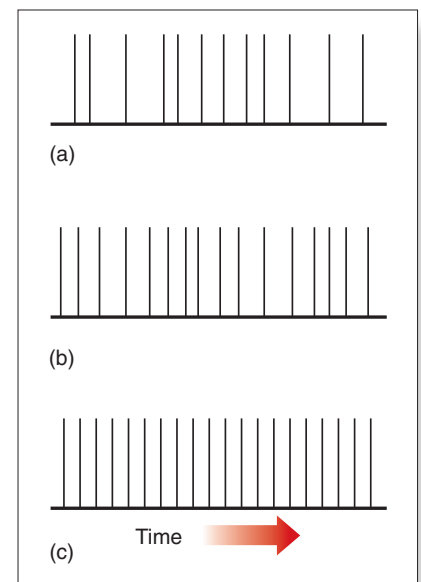
Although all of these discoveries about the nature of neurons and the signals that travel in them were extremely important (and garnered a number of Nobel Prizes for their discoverers), our main interest is not in how axons transmit signals, but in how these signals contribute to the operation of the mind. So far, our description of how signals are transmitted is analogous to describing how the Internet transmits electrical signals, without describing how the signals are transformed into words and pictures that people can understand. Adrian was acutely aware that it was important to go beyond simply describing nerve signals, so he did a series of experiments to relate nerve signals to stimuli in the environment and therefore to people's experience.

Adrian studied the relation between nerve firing and sensory experience by measuring how the firing of a neuron from a receptor in the skin changed as he applied more pressure to the skin. What he found was that the shape and height of the action potential remained the same as he increased the pressure, but the **rate of nerve firing**—that is, the number of action potentials that traveled down the axon per second—increased (**Figure 2.7**). From this result, Adrian drew a connection between nerve firing and experience. He describes this connection in his book *The Basis of Sensation* (1928) by stating that if nerve impulses “are crowded closely together the sensation is intense, if they are separated by long intervals the sensation is correspondingly feeble” (p. 7).

What Adrian is saying is that electrical signals are *representing* the intensity of the stimulus, so pressure that generates “crowded” electrical signals feels stronger than pressure that



► **Figure 2.6** (a) A series of action potentials displayed on a time scale that makes each action potential appear as a thin line. (b) Changing the time scale reveals the shape of one of the action potentials.



► **Figure 2.7** Action potentials recorded from an axon in response to three levels of pressure stimulation on the skin: (a) light, (b) medium, and (c) strong. Increasing stimulus intensity causes an increase in the rate of nerve firing.

generates signals separated by long intervals. Later experiments demonstrated similar results for vision. Presenting high-intensity light generates a high rate of nerve firing and the light appears bright; presenting lower intensity light generates a lower rate of nerve firing and the light appears dimmer. Thus, the rate of neural firing is related to the intensity of stimulation, which, in turn, is related to the magnitude of an experience, such as feeling pressure on the skin or experiencing the brightness of a light.

Going beyond Adrian's idea that the magnitude of experience is related to the rate of nerve firing, we can ask, how is the *quality* of experience represented in neural firing? For the senses, quality *across the senses* refers to the different experience associated with each of the senses—perceiving light for vision, sound for hearing, smells for olfaction, and so on. We can also ask about quality *within a particular sense*, such as for vision: color, movement, an object's shape, or the identity of a person's face.

One way to answer the question of how action potentials determine different qualities is to propose that the action potentials for each quality might look different. However, Adrian ruled out that possibility by determining that all action potentials have basically the same height and shape. If all nerve impulses are basically the same whether they are caused by seeing a red fire engine or remembering what you did last week, how can these impulses stand for different qualities? The short answer to this question is that different qualities of stimuli, and also different aspects of experience, activate different neurons and areas in the brain. We begin the long answer to this question in the next section by taking up the idea of representation, which we introduced in Chapter 1.

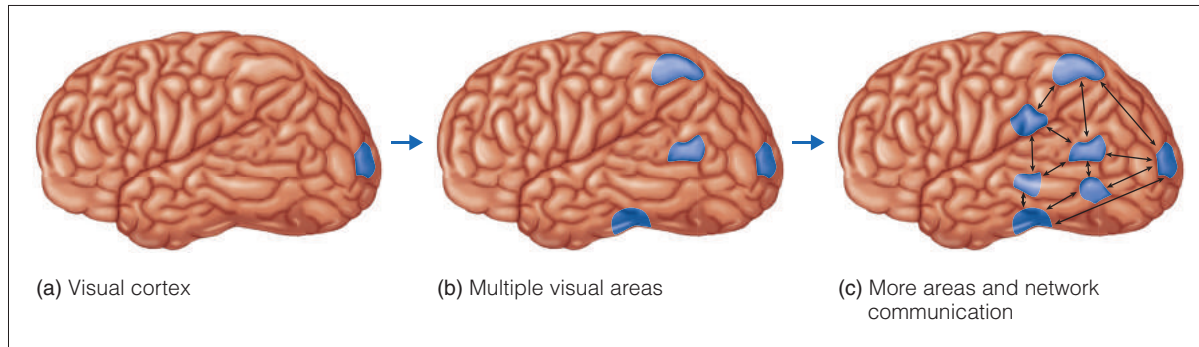
► Representation by Neural Firing

In Chapter 1, we defined the mind as *a system that creates representations of the world so that we can act within it to achieve our goals* (page 6). The key word in this definition is *representations*, because what it means is that everything we experience is the result of something that *stands for* that experience. Putting this in neural terms, the principle of neural representation states that everything a person experiences is based on representations in the person's nervous system. Adrian's pioneering research on how nerve impulses represent the intensity of a stimulus, in which he related high nerve firing to feeling greater pressure, marks the beginning of research on neural representation. We now move ahead to the 1960s to describe early research that involved recording from single neurons in the brain.

The Story of Neural Representation and Cognition: A Preview

In the 1960s, researchers began focusing on recording from single neurons in the primary visual receiving area, the place where signals from the eye first reach the cortex (Figure 2.8a). The question being asked in these experiments was “what makes this neuron fire?” Vision dominated early research because stimuli could be easily controlled by creating patterns of light and dark on a screen and because a lot was already known about vision.

But as research progressed, researchers began recording from neurons in areas outside the primary visual area and discovered two key facts: (1) Many neurons at higher levels of the visual system fire to complex stimuli like geometrical patterns and faces; and (2) a specific stimulus causes neural firing that is distributed across many areas of the cortex (Figure 2.8b). Vision, it turns out, isn't created only in the primary visual receiving area, but in many different areas. Later research, extending beyond vision, found similar results for other cognitions. For example, it was discovered that memory is not determined by a single “memory area,” because there are a number of areas involved in creating memories and remembering them later. In short, it became obvious that large areas of the brain are involved in creating cognition.



► **Figure 2.8** (a) Early work on neural representation and cognition focused on recording from single neurons in the visual cortex, where signals first arrive at the cortex. (b) Researchers then began to explore other places in the brain and found that visual stimulation causes activity that is distributed across many areas of the cortex. (c) Recent work has focused on looking at how these distributed areas are connected by neural networks and how activity flows in these networks. Note that, with the exception of the visual area in (a), the locations of the areas in this figure do not represent the locations of actual areas. They are for illustrative purposes only.

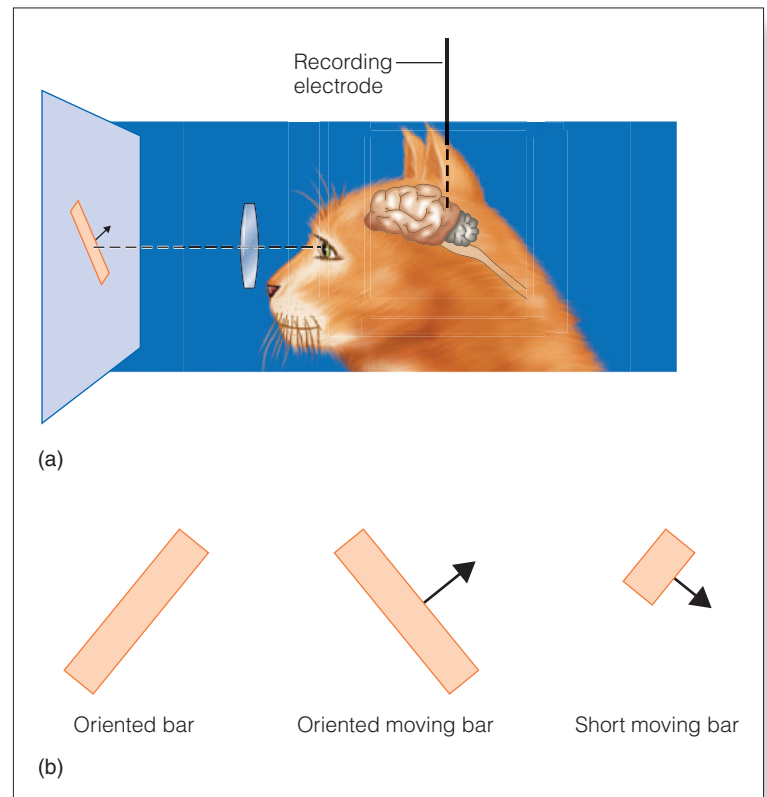
As it became clear that understanding neural representation involves casting a wide net across the brain, many researchers began considering the way different areas are connected to one another. The idea of neural signals transmitted between many destinations in an interconnected brain has led to today's conception of the brain as containing a vast highway system that can be described in terms of "neural networks" (Figure 2.8c). We will now fill in the details, beginning with the discovery of neural feature detectors.

Feature Detectors

One possible answer to the question "how can nerve impulses stand for different qualities?" is that perhaps there are neurons that fire only to specific qualities of stimuli. Early research found some evidence for this (Hartline, 1940; Kuffler, 1953), but the idea of neurons that respond to specific qualities was brought to the forefront by a series of papers by David Hubel and Thorsten Wiesel, which would win them the Nobel Prize in 1981.

In the 1960s, Hubel and Wiesel started a series of experiments in which they presented visual stimuli to cats, as shown in Figure 2.9a, and determined which stimuli caused specific neurons to fire. They found that each neuron in the visual area of the cortex responded to a specific type of stimulation presented to a small area of the retina. Figure 2.9b shows some of the stimuli that caused neurons in and near the visual cortex to fire (Hubel, 1982; Hubel & Wiesel, 1959, 1961, 1965). They called these neurons **feature detectors** because they responded to specific stimulus features such as orientation, movement, and length.

The idea that feature detectors are linked to perception was supported by many different experiments. One of these experiments involved a phenomenon called



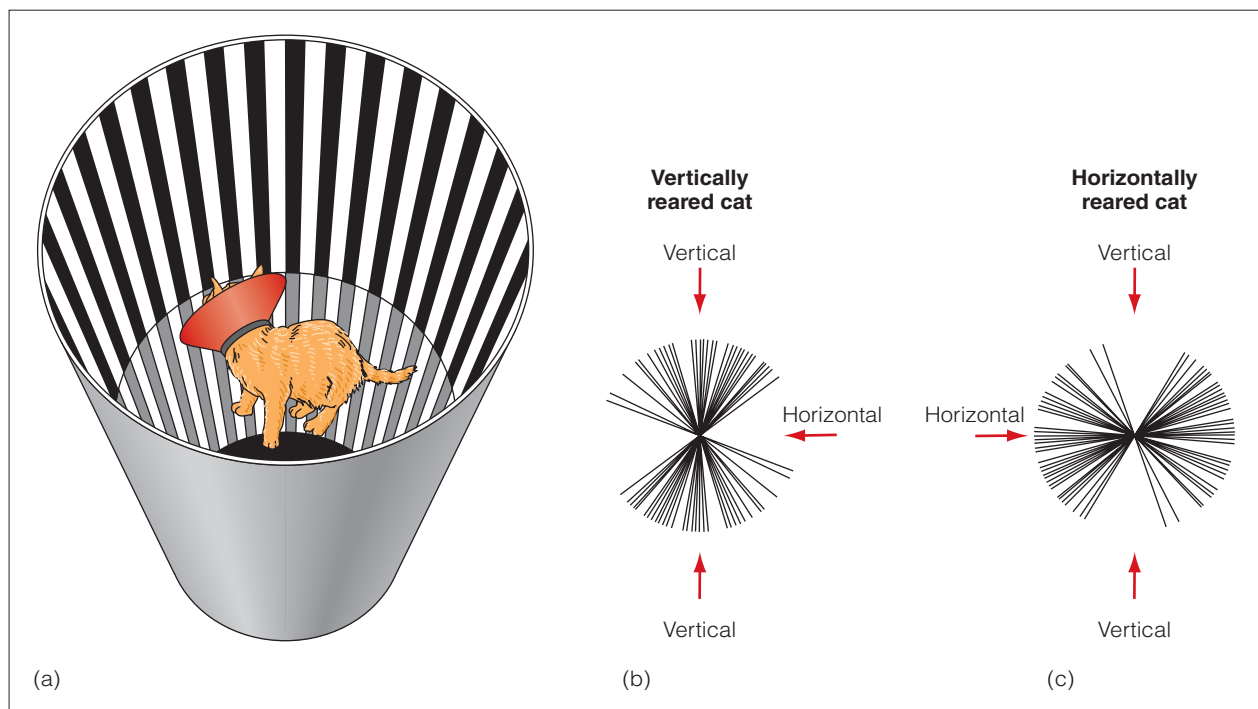
► **Figure 2.9** (a) An experiment in which electrical signals are recorded from the visual system of an anesthetized cat that is viewing stimuli presented on the screen. The lens in front of the cat's eye ensures that the images on the screen are focused on the cat's retina. The recording electrode is not shown. (b) A few of the types of stimuli that cause neurons in the cat's visual cortex to fire.

experience-dependent plasticity, in which the structure of the brain is changed by experience. For example, when a kitten is born, its visual cortex contains feature detectors that respond to oriented bars (see Figure 2.9). Normally, the kitten's visual cortex contains neurons that respond to all orientations, ranging from horizontal to slanted to vertical, and when the kitten grows up into a cat, the cat has neurons that can respond to all orientations.

But what would happen if kittens were reared in an environment consisting only of verticals? Colin Blakemore and Graham Cooper (1970) answered this question by rearing kittens in a space in which they saw only vertical black and white stripes on the walls (Figure 2.10a). After being reared in this vertical environment, kittens batted at a moving vertical stick but ignored horizontal objects. The basis of this lack of response to horizontals became clear when recording from neurons in the kittens' brains revealed that the visual cortex had been reshaped so it contained neurons that responded mainly to verticals and had no neurons that responded to horizontals (Figure 2.10b). Similarly, kittens reared in an environment consisting only of horizontals ended up with a visual cortex that contained neurons that responded mainly to horizontals (Figure 2.10c). Thus, the kittens' brains had been shaped to respond best to the environment to which they had been exposed.

Blakemore and Cooper's experiment is important because it is an early demonstration of experience-dependent plasticity. Their result also has an important message about neural representation: When a kitten's cortex contained mainly vertically sensitive neurons, the kitten perceived only verticals, and a similar result occurred for horizontals. This result supports the idea that perception is determined by neurons that fire to specific qualities of a stimulus (orientation, in this case).


This knowledge that neurons in the visual system fire to specific types of stimuli led to the idea that each of the thousands of neurons that fire when we look at a tree fire to different features of the tree. Some neurons fire to the vertically oriented trunk, others to



► **Figure 2.10** (a) Striped tube used in Blakemore and Cooper's (1970) selective rearing experiments. (b) Distribution of orientations that caused maximum firing for 72 cells from a cat reared in an environment of vertical stripes and (c) for 52 cells from a cat reared in an environment of horizontal stripes.

the variously oriented branches, and some to more complex combinations of a number of features. **The idea that the tree is represented by the combined response of many feature detectors is similar to building objects by combining building blocks like Legos.** But it is important to realize that the visual cortex is an early stage of visual processing, and that vision depends on signals that are sent from the visual cortex to other areas of the brain.

Figure 2.11 indicates the location of the **visual cortex** in the human brain, as well as additional areas that are involved in vision, and some other areas we will be discussing later. The vision areas are part of a vast network of areas that make up about 30 percent of the cortex (Felleman & Van Essen, 1991). Some of these visual areas receive signals directly from the visual cortex. Others are part of a sequence of interconnected neurons, some of which are far down the line from the visual cortex. Following Hubel and Wiesel's pioneering research, other researchers who began exploring these "higher" levels of the visual pathway discovered neurons that respond to stimuli more complex than oriented lines.




► **Figure 2.11** Some of the structures of the human brain that we will be referring to in this chapter. Pointers indicate the locations of these areas, each of which extends over an area of the cortex.

Neurons That Respond to Complex Stimuli

How are complex stimuli represented by the firing of neurons in the brain? One answer to this question began to emerge in the laboratory of Charles Gross. Gross's experiments, in which he recorded from single neurons in the monkey's **temporal lobe** (Figure 2.11), required a great deal of endurance by the researchers, because the experiments typically lasted 3 or 4 days. In these experiments, the results of which were reported in now classic papers in 1969 and 1972 (Gross et al., 1969, 1972), Gross's research team presented a variety of different stimuli to anesthetized monkeys. On a projection screen like the one shown in Figure 2.9a, they presented lines, squares, and circles. Some stimuli were light and some dark.

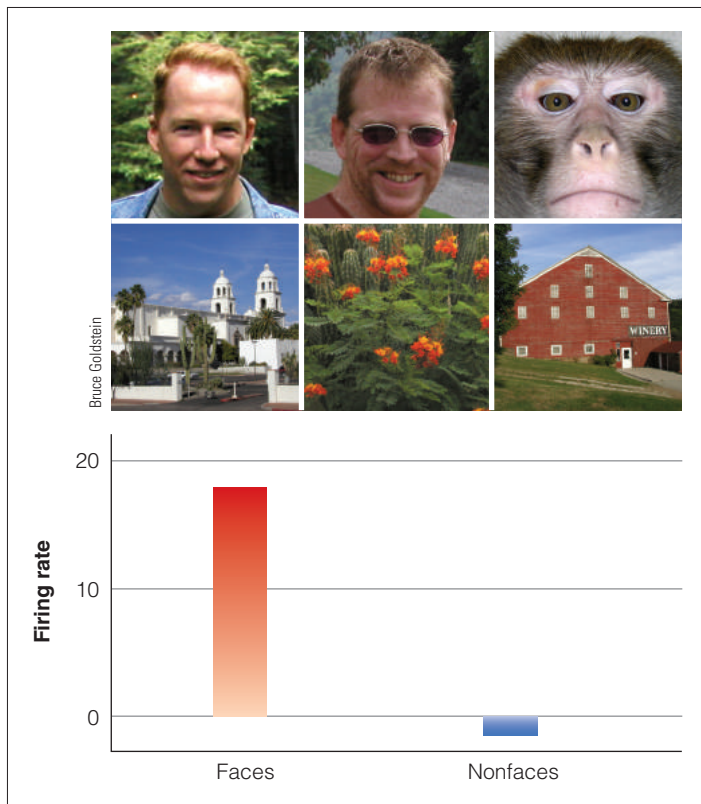
The discovery that neurons in the temporal lobe respond to complex stimuli came a few days into one of their experiments, when they had found a neuron that refused to respond to any of the standard stimuli, like oriented lines or circles or squares. Nothing worked, until one of the experimenters pointed at something in the room, casting a shadow of his hand on the screen. **When this hand shadow caused a burst of firing, the experimenters knew they were on to something and began testing the neuron with a variety of stimuli, including cutouts of a monkey's hand.** After a great deal of testing, they determined that this neuron responded best to a handlike shape with fingers pointing up (far-right stimuli in **Figure 2.12**) (Rocha-Miranda, 2011; also see Gross, 2002). After expanding the types of stimuli presented, they also found some neurons that responded best to faces. Later researchers extended these results and provided many examples of neurons that respond to faces but don't respond to other types of stimuli (Perrett et al., 1982; Rolls, 1981) (**Figure 2.13**).

Let's stop for a moment and consider the results we have presented so far. We saw that neurons in the visual cortex respond to simple stimuli like oriented bars, neurons in the temporal lobe



► **Figure 2.12** Some of the shapes used by Gross et al. (1972) to study the responses of neurons in the temporal lobe of the monkey's cortex. The shapes are arranged in order of their ability to cause the neuron to fire, from none (1) to little (2 and 3) to maximum (6).

(Source: Based on Gross et al., 1972.)



► **Figure 2.13** Firing rate, in nerve impulses per second, of a neuron in the monkey's temporal lobe that responds to face stimuli but not to nonface stimuli.

(Source: Based on E. T. Rolls & M. J. Tovee, 1995.)

to three different faces. Only neuron 4 responds to Bill's face, only neuron 9 responds to Mary's face, and only neuron 6 responds to Raphael's face. Also note that the neuron specialized to respond only to Bill, which we can call a "Bill neuron," does not respond to Mary or Raphael. In addition, other faces or types of objects would not affect this neuron. It fires only to Bill's face.

Although the idea of specificity coding is straightforward, it is unlikely to be correct. Even though there are neurons that respond to faces, these neurons usually respond to a number of different faces (not just Bill's). There are just too many different faces and other objects (and colors, tastes, smells, and sounds) in the world to have a separate neuron dedicated to each object. An alternative to the idea of specificity coding is that a number of neurons are involved in representing an object.

Population coding is the representation of a particular object by the pattern of firing of a large number of neurons (Figure 2.14b). According to this idea, Bill's, Mary's and Raphael's faces are each represented by a different pattern. An advantage of population coding is that a large number of stimuli can be represented, because large groups of neurons can create a huge number of different patterns. There is good evidence for population coding in the senses and for other cognitive functions as well. But for some functions, a large number of neurons aren't necessary.

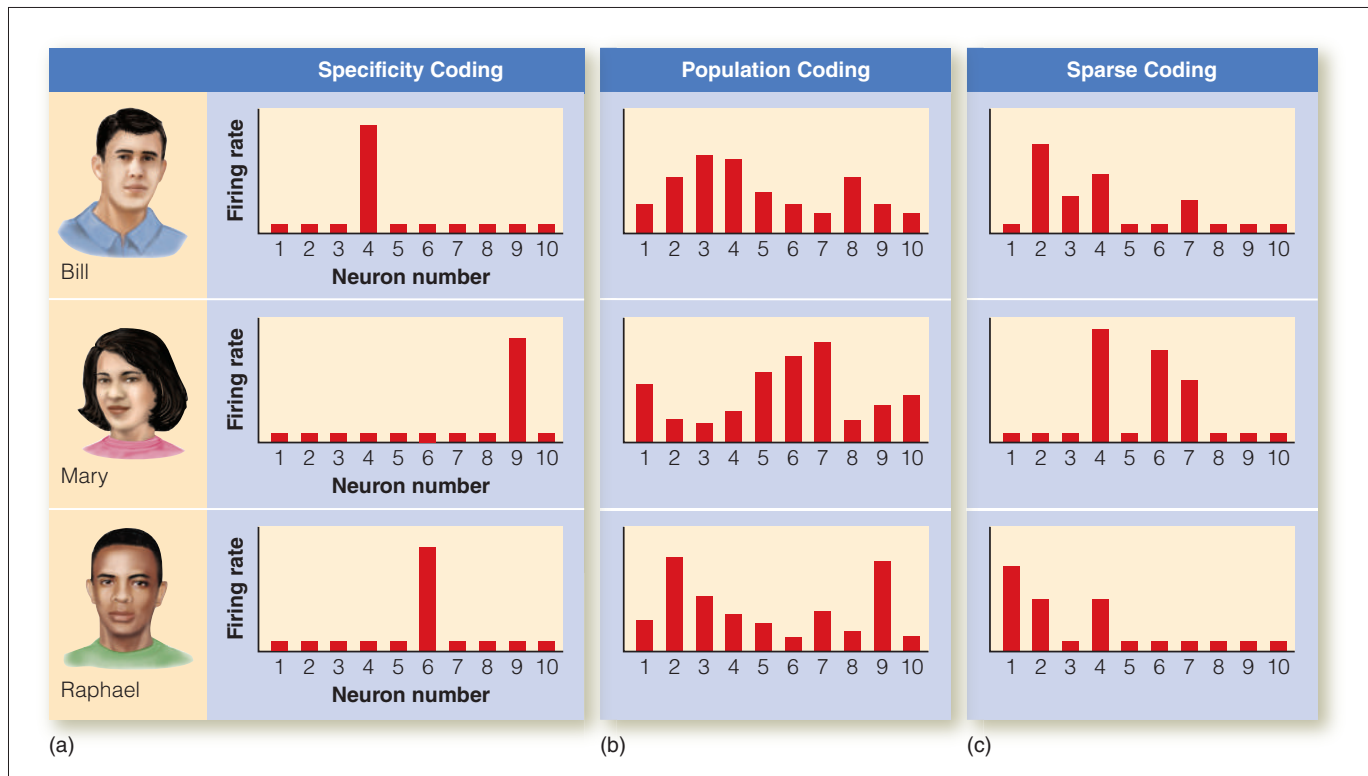
Sparse coding occurs when a particular object is represented by a pattern of firing of only a small group of neurons, with the majority of neurons remaining silent. As shown in Figure 2.14c, sparse coding would represent Bill's face by the pattern of firing of a few neurons (neurons 2, 3, 4, and 7). Mary's face would be signaled by the pattern of firing of a few different neurons (neurons 4, 6, and 7), but possibly with some overlap with the neurons representing Bill, and Raphael's face would have yet another pattern (neurons 1, 2, and 4).

respond to complex geometrical stimuli, and neurons in another area of the temporal lobe respond to faces. What is happening is that neurons in the visual cortex that respond to relatively simple stimuli send their axons to higher levels of the visual system, where signals from many neurons combine and interact; neurons at this higher level, which respond to more complex stimuli such as geometrical objects, then send signals to even higher areas, combining and interacting further and creating neurons that respond to even more complex stimuli such as faces. This progression from lower to higher areas of the brain is called **hierarchical processing**.

Does hierarchical processing solve the problem of neural representation? Could it be that higher areas of the visual system contain neurons that are specialized to respond only to a specific object, so that object would be represented by the firing of that one type of specialized neuron? As we will see, this is probably not the case, because neural representation most likely involves a number of neurons working together.

Sensory Coding

The problem of neural representation for the senses has been called the **problem of sensory coding**, where the **sensory code** refers to how neurons represent various characteristics of the environment. The idea that an object could be represented by the firing of a specialized neuron that responds only to that object is called **specificity coding**. This is illustrated in Figure 2.14a, which shows how a number of neurons respond

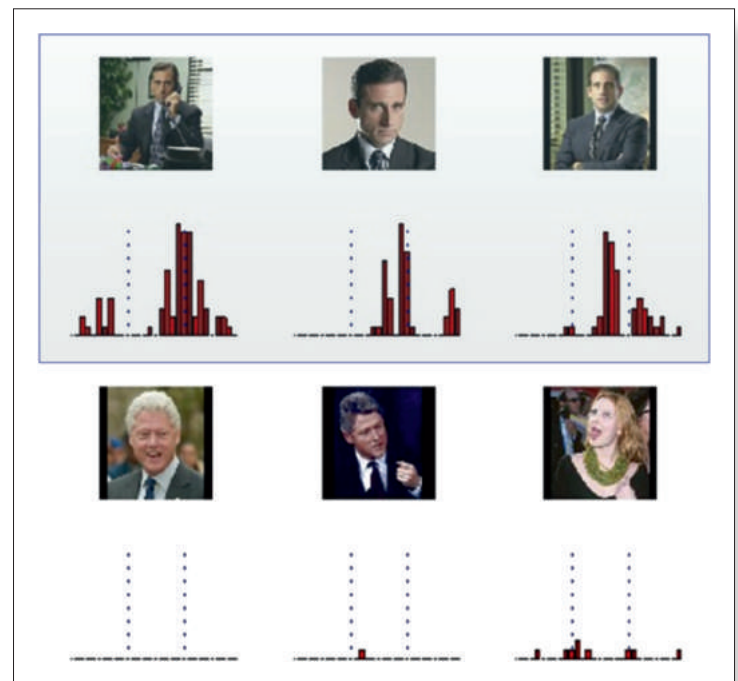


► **Figure 2.14** Three types of coding: (a) Specificity coding. The response of 10 different neurons to each face on the left is shown. Each face causes a different neuron to fire. (b) Population coding. The face's identity is indicated by the pattern of firing of a large number of neurons. (c) Sparse coding. The face's identity is indicated by the pattern of firing of a small group of neurons.

Notice that a particular neuron can respond to more than one stimulus. For example, neuron 4 responds to all three faces, although most strongly to Mary's.

Recently, neurons were discovered when recording from the temporal lobe of patients undergoing brain surgery for epilepsy. (Stimulating and recording from neurons is a common procedure before and during brain surgery, because it makes it possible to determine the exact layout of a particular person's brain.) **These neurons responded to very specific stimuli. Figure 2.15 shows the records for a neuron that responded to pictures of the actor Steve Carell and not to other people's faces (Quiroga et al., 2007).** However, the researchers who discovered this neuron (as well as other neurons that responded to other people) point out that they had only 30 minutes to record from these neurons, and that if more time were available, it is likely that they would have found other faces that would cause this neuron to fire. Given the likelihood that even these special neurons are likely to fire to more than one stimulus, Quiroga and coworkers (2008) suggested that **their neurons are probably an example of sparse coding.**

There is also other evidence that the code for representing objects in the visual system, tones in the auditory system, and odors in the olfactory system may involve the pattern of



► **Figure 2.15** Records from a neuron in the temporal lobe that responded to different views of Steve Carell (top records) but did not respond to pictures of other well-known people (bottom records). (Source: Quiroga et al., 2008)

activity across a relatively small number of neurons, as sparse coding suggests (Olshausen & Field, 2004).

Memories are also represented by the firing of neurons, but **there is a difference between representation of perceptions and representation of memories.** The neural firing associated with experiencing a perception is associated with what is happening as a stimulus is present. Firing associated with memory is associated with information about the past that **has been stored in the brain.** We know less about the actual form of this stored information for memory, but it is likely that the basic principles of population and sparse coding also operate for memory, with specific memories being represented by particular patterns of stored information that result in a particular pattern of nerve firing when we experience the memory.

Saying that individual neurons and groups of neurons contain information for perception, memory, and other cognitive functions is the first step toward understanding representation. The next step involves looking at organization: how different types of neurons and functions are organized within the brain.

TEST YOURSELF 2.1

1. Describe the idea of levels of analysis.
2. How did early brain researchers describe the brain in terms of a nerve net? How does the idea of individual neurons differ from the idea of a nerve net?
3. Describe the research that led Cajal to propose the neuron doctrine.
4. Describe the structure of a neuron. Describe the synapse and neural circuits.
5. How are action potentials recorded from a neuron? What do these signals look like, and what is the relation between action potentials and stimulus intensity?
6. How has the question of how different perceptions can be represented by neurons been answered? Consider both research involving recording from single neurons and ideas about sensory coding.
7. How is neural representation for memory different from neural representation for perception? How is it similar?

► Localized Representation

One of the basic principles of brain organization is **localization of function**—specific functions are served by specific areas of the brain. Many cognitive functions are served by the **cerebral cortex**, which is a layer of tissue about 3 mm thick that covers the brain (Fischl & Dale, 2000). The cortex is the wrinkled covering you see when you look at an intact brain (Figure 2.11). Other functions are served by **subcortical areas** that are located below the cortex. Early evidence for localization of function came from **neuropsychology**—the study of the behavior of people with brain damage.

Localization Determined by Neuropsychology

In the early 1800s, an accepted principle of brain function was **cortical equipotentiality**, the idea that **the brain operated as an indivisible whole as opposed to specialized areas** (Flourens, 1824; Pearce, 2009). But in 1861, Paul Broca published work based on his study of patients who had suffered brain damage due to strokes that caused disruption of the blood supply to the brain. These strokes caused damage to an area in the frontal lobe that came to be called **Broca's area** (Figure 2.16).

One of Broca's patients was famously called *Tan*, because the word *tan* was the only word he could say. Other patients with frontal lobe damage could say more, but their speech

was slow and labored and often had jumbled sentence structure. Here is an example of the speech of a modern patient, who is attempting to describe when he had his stroke, which occurred when he was in a hot tub:

Alright. . . . Uh ... stroke and un. . . . I . . . huh tawanna guy. . . .
H . . . h . . . hot tub and. . . . And the. . . . Two days when uh. . . .
Hos . . . uh. . . . Huh hospital and uh . . . amet . . . am . . . ambulance.
(Dick et al., 2001, p. 760)

Patients with this problem—**slow, labored, ungrammatical speech caused by damage to Broca's area**—are diagnosed as having **Broca's aphasia**. The fact that damage to a specific area of the brain caused a specific deficit of behavior was striking evidence against the idea of equipotentiality and for the idea of localization of function.

Eighteen years after Broca reported on his frontal lobe patients, Carl Wernicke (1879) described a number of patients who had damage to an area in their temporal lobe that came to be called **Wernicke's area**. **Wernicke's patients produced speech that was fluent and grammatically correct but tended to be incoherent**. Here is a modern example of the speech of a patient with **Wernicke's aphasia**:

It just suddenly had a feffort and all the feffort had gone with it. It even stepped my horn. They took them from earth you know. They make my favorite nine to severed and now I'm a been habed by the uh stam of fortment of my annulment which is now forever. (Dick et al., 2001, p. 761)

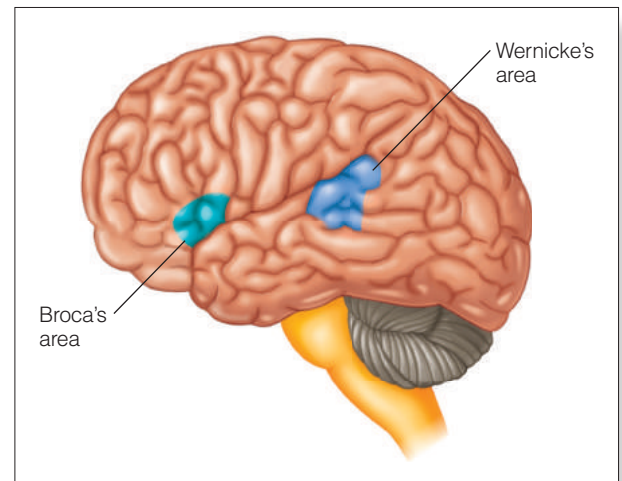
Patients such as this not only produce meaningless speech but are unable to understand other people's speech. Their primary problem is their inability to match words with their meanings, with the defining characteristic of Wernicke's aphasia being the absence of normal grammar (Traxler, 2012).

Broca's and Wernicke's observations showed that different aspects of language—production of language and comprehension of language—were served by different areas in the brain. As we will see later in this chapter, modern research has shown that the strict separation of language functions in different areas was an oversimplification. Nonetheless, Broca's and Wernicke's 19th-century observations set the stage for later research that confirmed the idea of localization of function.

Further evidence for localization of function came from studies of the effect of brain injury in wartime. Studies of Japanese soldiers in the Russo-Japanese war of 1904–1905 and Allied soldiers in World War I showed that **damage to the occipital lobe of the brain, where the visual cortex is located (Figure 2.11), resulted in blindness, and that there was a connection between the area of the occipital lobe that was damaged and the place in visual space where the person was blind (Glickstein & Whitteridge, 1987; Holmes & Lister, 1916; Lanska, 2009)**. For example, damage to the left part of the occipital lobe caused an area of blindness in the upper-right part of visual space.

As noted earlier, other areas of the brain have also been associated with specific functions. **The auditory cortex, which receives signals from the ears, is in the upper temporal lobe and is responsible for hearing. The somatosensory cortex, which receives signals from the skin, is in the parietal lobe and is responsible for perceptions of touch, pressure, and pain.** The **frontal lobe** receives signals from all of the senses and is responsible for coordination of the senses, as well as higher cognitive functions like thinking and problem solving.

Another effect of brain damage on visual functioning, reported in patients who have damage to the temporal lobe on the lower-right side of the brain, is **prosopagnosia—an inability to recognize faces**. People with prosopagnosia can tell that a face is a face, but they can't recognize whose face it is, even for people they know well such as friends and family



► **Figure 2.16** Broca's area, in the frontal lobe, and Wernicke's area, in the temporal lobe, were identified in early research as being specialized for language production and comprehension, respectively.

members. In some cases, people with prosopagnosia look into a mirror and, seeing their own image, wonder who the stranger is looking back at them (Burton et al., 1991; Hecaen & Angelergues, 1962; Parkin, 1996).

One of the goals of the neuropsychology research we have been describing is to determine whether a particular area of the brain is specialized to serve a particular cognitive function. Although it might be tempting to conclude, based on a single case of prosopagnosia, that the damaged brain area in the lower temporal lobe is responsible for recognizing faces, modern researchers realized that to reach more definite conclusions about the function of a particular area, it is necessary to test a number of different patients with damage to different brain areas in order to demonstrate a *double dissociation*.

METHOD Demonstrating a Double Dissociation

A **double dissociation** occurs if damage to one area of the brain causes function A to be absent while function B is present, and damage to another area causes function B to be absent while function A is present. To demonstrate a double dissociation, it is necessary to find two people with brain damage that satisfy the above conditions.

Double dissociations have been demonstrated for face recognition and object recognition, by finding patients with brain damage who can't recognize faces (Function A) but who can recognize objects (Function B), and other patients, with brain damage in a different area, who can't recognize objects (Function B) but who can recognize faces (Function A) (McNeal & Warrington, 1993; Moscovitch et al., 1997). The importance of demonstrating a double dissociation is that it enables us to conclude that functions A and B are served by different mechanisms, which operate independently of one another.

The results of the neuropsychology studies described above indicate that face recognition is served by one area in the temporal lobe and that this function is separate from mechanisms associated with recognizing other types of objects, which is served by another area of the temporal lobe. Neuropsychological research has also identified areas that are important for perceiving motion and for different functions of memory, thinking, and language, as we will see later in this book.

Localization Determined by Recording from Neurons

Another tool for demonstrating localization of function is recording from single neurons. Numerous studies, mostly on animals, used single-neuron recording to demonstrate localization of function. For example, Doris Tsao and coworkers (2006) found that 97 percent of neurons within a small area in the lower part of a monkey's temporal lobe responded to pictures of faces but not to pictures of other types of objects. This "face area," as it turns out, is located near the area in humans that is associated with prosopagnosia. The idea that our perception of faces is associated with a specific area of the brain is also supported by research using the technique of **brain imaging** (see Chapter 1, page 18), which makes it possible to determine which areas of the brains of humans are activated by different cognitions.

Localization Demonstrated by Brain Imaging

We noted in Chapter 1 that technological advances that cause a shift in the way science is done can be called a revolution. On that basis, it could be argued that the introduction of the brain-scanning techniques positron emission tomography (PET) in 1976 and functional magnetic resonance imaging (fMRI) in 1990 marked the beginning of the "imaging revolution."

As you will see throughout this book, brain scanning, and especially fMRI, has played an important role in understanding the physiological basis of cognition. Here we consider what fMRI research tells us about localization of function in the brain. We begin by describing the basic principle behind fMRI.

METHOD Brain Imaging

Functional magnetic resonance imaging (fMRI) takes advantage of the fact that neural activity causes the brain to bring in more oxygen, which binds to hemoglobin molecules in the blood. This added oxygen increases the magnetic properties of the hemoglobin, so when a magnetic field is presented to the brain, these more highly oxygenated hemoglobin molecules respond more strongly to the magnetic field and cause an increase in the fMRI signal.

The setup for an fMRI experiment is shown in **Figure 2.17a**, with the person's head in the scanner. As a person engages in a cognitive task such as perceiving an image, the activity of the brain is determined. Activity is recorded in **voxels**, which are small, cube-shaped areas of the brain about 2 or 3 mm on a side. Voxels are not brain structures but are simply small units of analysis created by the fMRI scanner. One way to think about voxels is that they are like the small, square pixels that make up digital photographs or the images on your computer screen, but because the brain is three-dimensional, voxels are small cubes rather than small squares. **Figure 2.17b** shows the result of an fMRI scan. Increases or decreases in brain activity associated with cognitive activity are indicated by colors, with specific colors indicating the amount of activation.

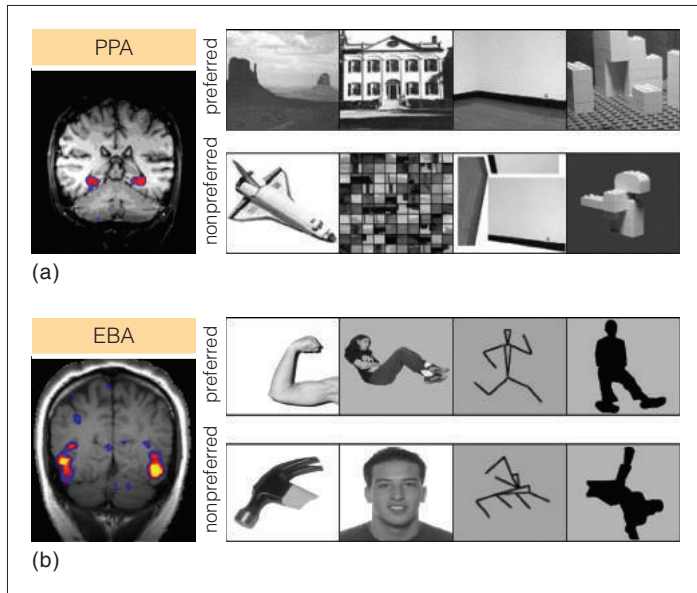
It bears emphasizing that these colored areas do not appear as the brain is being scanned. They are determined by a procedure that involves taking into account how the brain is responding when the person is not engaged in a task and the change in activity triggered by the task. Complex statistical procedures are used to determine the **task-related fMRI**—the change in brain activity that can be linked specifically to the task. The results of these calculations for each voxel are then displayed as colorful activation patterns, like the one in **Figure 2.17b**.

Many of the brain-imaging experiments that have provided evidence for localization of function have involved determining which brain areas were activated when people observed pictures of different objects.



► **Figure 2.17** (a) Person in a brain scanner. (b) fMRI record. Colors indicate locations of increases and decreases in brain activity. Red and yellow indicate increases in brain activity; blue and green indicate decreases.

(Source: Part b from Ishai et al., 2000)



► **Figure 2.18** (a) The parahippocampal place area (PPA) is activated by places (top row) but not by other stimuli (bottom row). (b) The extrastriate body area (EBA) is activated by bodies (top) but not by other stimuli (bottom). (Source: Chalupa & Werner, 2003)

Looking at Pictures We've already seen how neuropsychology research and single neuron recording identified areas that are involved in perceiving faces. A face area has also been identified by having people in a brain scanner look at pictures of faces. This area, which is called the **fusiform face area (FFA)** because it is in the fusiform gyrus on the underside of the temporal lobe (Kanwisher et al., 1997), is the same part of the brain that is damaged in cases of prosopagnosia (Figure 2.11).


Further evidence for localization of function comes from fMRI experiments that have shown that perceiving pictures representing indoor and outdoor scenes like those shown in **Figure 2.18a** activates the **parahippocampal place area (PPA)** (Aguirre et al., 1998; Epstein et al., 1999). Apparently, what is important for this area is information about spatial layout, because increased activation occurs when viewing pictures both of empty rooms and of rooms that are completely furnished (Kanwisher, 2003). Another specialized area, the **extrastriate body area (EBA)**, is activated by pictures of bodies and parts of bodies (but not by faces), as shown in **Figure 2.18b** (Downing et al., 2001).

Looking at Movies Our usual experience, in everyday life, involves seeing scenes that contain many different objects, some of which are moving. Therefore, Alex Huth and coworkers (2012) conducted an fMRI experiment using stimuli similar to what we see in the environment, by having participants view film clips. Huth's participants viewed 2 hours of film clips while in a brain scanner. To analyze how the voxels in these participants' brains responded to different objects and actions in the films, Huth created a list of 1,705 different objects and action categories and determined which categories were present in each film scene.

Figure 2.19 shows four scenes and the categories (labels) associated with them. By determining how each voxel responded to each scene and then analyzing his results using a complex statistical procedure, Huth was able to determine what kinds of stimuli each voxel

► **Figure 2.19** Four frames from the movies viewed by participants in Huth et al.'s (2012) experiment. The words on the right indicate categories that appear in the frames (n = noun, v = verb). (Huth et al., 2012)

Movie Clip	Labels	Movie Clip	Labels
	butte.n desert.n sky.n cloud.n brush.n		city.n expressway.n skyscraper.n traffic.n sky.n
	woman.n talk.v gesticulate.v book.n		bison.n walk.v grass.n stream.n



► **Figure 2.20** The results of Huth et al.'s (2012) experiment, showing locations on the brain where the indicated categories are most likely to activate the brain. (Source: Courtesy of Alex Huth)

responded to. For example, one voxel responded well when streets, buildings, roads, interiors, and vehicles were present.

Figure 2.20 shows the types of stimuli that cause voxels across the surface of the brain to respond. **Objects and actions similar to each other are located near each other in the brain.** The reason there are two areas for humans and two for animals is that each area represents **different features related to humans or animals.** For example, the area labeled “Humans” at the bottom of the brain (which is actually on the underside of the brain) corresponds to the fusiform face area (Figure 2.11), which responds to all aspects of faces. The area labeled “Humans” higher on the brain responds specifically to facial expressions. The areas labeled “Talking” correspond to Broca’s and Wernicke’s areas.

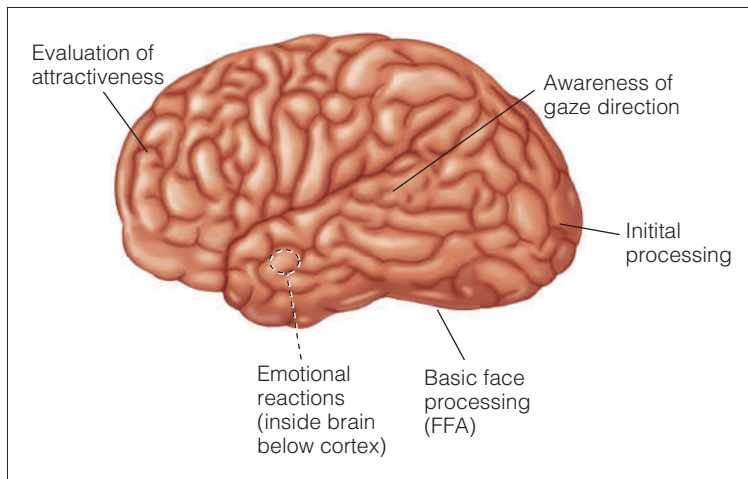
The results in Figure 2.20 present an interesting paradox. On one hand, the results confirm the earlier research that identified specific areas of the brain responsible for the perception of specific types of stimuli like faces, places, and bodies. On the other hand, these new results reveal a map that stretches over a large area of the cortex. As we will now see, even though there is a great deal of evidence for localization of function, we need to consider the brain as a whole in order to understand the physiological basis of cognition.

► Distributed Representation

Let’s consider Huth’s map of categories in the brain in Figure 2.20, which shows that there are two locations for “Humans.” The explanation—that different areas respond to different features of humans—illustrates a central principle of cognition: **most of our experience is multidimensional.** That is, even simple experiences involve combinations of different qualities. Consider, for example, looking at a person’s face.


Looking at a Face

Looking at a face triggers responses to many different aspects of the face. So in addition to identifying an object as a face (“that’s a face”), we also respond to the following additional aspects of faces: (1) emotional aspects (“she is smiling, so she is probably happy,” “looking at his face makes me happy”); (2) where someone is looking (“she’s looking at me”);



► **Figure 2.21** Areas of the brain that are activated by different aspects of faces.

(Sources: Adapted from Ishai, 2008; based on data from Calder et al., 2007; Gobbini & Haxby, 2007; Grill-Spector et al., 2004; Haxby et al., 2000; Ishai et al., 2004.)



► **Figure 2.22** Brain showing areas activated by episodic and semantic memories. Yellow = episodic. Blue = semantic.

(Source: From Levine et al., 2004)

(3) how parts of the face move (“I can understand him better by watching his lips move”); (4) how attractive a face is (“he has a handsome face”); and (5) whether the face is familiar (“I remember her from somewhere”). This multidimensional response to faces is reflected in distributed neural responses throughout the cortex (**Figure 2.21**).

The fact that looking at a face activates many areas of the brain is called **distributed representation**. Cognitions, be they perceptions from looking at something, or processes such as remembering or thinking, activate numerous, sometimes widely separated, areas of the brain. Let’s examine two additional examples of distributed neural representation.

Remembering

Memories are complicated. Some memories, called short-term memories, last fleetingly, for only about 10 to 15 seconds unless repeated over and over, as you might do to remember a phone number you forgot to store in your cell phone. Other memories are longer, as your memory for something you did last week or even years ago. As we will see in Chapter 5, there is evidence that short-term and long-term memories are served by different areas of the brain (Curtis & D’Esposito, 2003; Harrison & Tong, 2009).

But memories also differ in another way. **Episodic memories** are memories for events in a person’s life, like remembering what you did yesterday. **Semantic memories** are memories for facts, like knowing that the capital of California is Sacramento. **Figure 2.22** shows the results of a brain-scanning experiment, which indicates that thinking about episodic and semantic memories activates different areas of the brain (Levine et al., 2004).

We will see in Chapters 5 through 7 that some areas of the brain play important roles in forming new memories and retrieving old ones, but there is also evidence that **remembering activates areas throughout the brain**. Memories can be visual (picturing someplace you often visit), auditory (remembering a favorite song), or olfactory (smell triggering memories for a familiar place). Memories often have emotional components, both good and bad (thinking about someone you miss). Most memories are combinations of many of these components, each of which activates different areas of the brain. Memories, therefore, create a symphony of neural activity throughout the brain.

Producing and Understanding Language


When we told the story about Broca and Wernicke, we focused on how their descriptions of two areas of the brain—one for producing speech, the other for comprehending speech—provided the impetus for the idea of localized functions. But in telling that story, we left something out: In addition to proposing an area for speech comprehension, Wernicke also suggested that language goes beyond isolated regions to include connections between them, and to other areas as well (Ross, 2010).

As it turns out, Wernicke’s proposal of connectivity stayed in the background after it was proposed, in favor of the idea of localized brain function, and **it wasn’t until the 20th century that his ideas about connectivity became well known and other researchers showed that explaining the physiology of language involved more than just two separate, localized language areas** (Geschwind, 1964; Ross, 2010).

Modern researchers have shown that damage to areas outside of Broca's and Wernicke's areas can cause problems in producing and understanding language (Ross, 2010). There is also evidence that nonlanguage functions are associated with parts of Broca's area (Federenko et al., 2012) and that processing of sentence grammar occurs throughout the language system (Blank et al., 2016). Results such as this led to a much more complex picture of how language is processed.

Figure 2.23 shows a modern diagram of language pathways. In this diagram, the language system is organized into two sets of pathways: one (in blue), which is involved with processing sounds, production of speech, and saying words, and the other (in red), which is involved in understanding words. Both sets of pathways are also involved in understanding sentences (Gierhan, 2013). This diagram represents the results of "research in progress," because much remains to be learned about how language is represented in the brain. However, there is no question that the representation of language is distributed throughout many areas.

One thing that the examples of perceiving faces, remembering, and language have in common is that they involve experiences that activate many separated brain areas, and there is evidence that many of these areas are linked either by direct neural connections or by being part of a number of interconnected structures. This brings us to an important new way of understanding the physiology of cognition that involves *neural networks*.



► **Figure 2.23** Pathways that have been reported to be involved in language processing. This diagram is based on the results of a large number of studies. Specific functions have been associated with each of the pathways, but the overall picture can be summarized by noting that the pathways shown in blue are involved in processing sounds, producing speech, and saying words, and the pathways shown in red are involved in understanding words, with both sets of pathways being involved in understanding sentences.

(Source: Gierhan, 2013)

► Neural Networks

Neural networks are interconnected areas of the brain that can communicate with each other (Bassett & Sporns, 2017). The idea of neural networks is a logical extension of the idea of distributed processing, because it makes sense that if many areas are involved in a particular type of cognition, that they might be connected.


As we tell the story of how researchers are discovering the properties of neural networks, we will be introducing four principles:

1. There are complex structural pathways called *networks* that form the brain's information highway.
2. Within these structural pathways there are functional pathways that serve different functions.
3. These networks operate dynamically, mirroring the dynamic nature of cognition.
4. There is a resting state of brain activity, so parts of the brain are active all the time, even when there is no cognitive activity.

We begin by considering how neural networks have been described structurally.

Structural Connectivity

Structural connectivity is the brain's "wiring diagram" created by nerve axons that connect different brain areas. Early researchers determined these connections using classical neuroanatomical techniques in which slices of brain tissue were stained to highlight axons, which enabled them to see the neural pathways with a microscope. But recently, new techniques have been developed that make more extensive mapping of the brain's connections possible.



Courtesy of Fernando Calamante and Elsevier

► **Figure 2.24** The connectome. Nerve tracts in the human brain determined by track-weighted imaging.
(Source: Calamante et al., 2013)

One of these techniques, called **track-weighted imaging (TWI)**, is based on detection of how water diffuses along the length of nerve fibers. **Figure 2.24** shows nerve tracts determined by this technique (Calamante, 2013). New techniques like this are constantly being developed to determine more precisely how areas of the brain communicate (Bressler & Menon, 2010; Sporns, 2015).

Pictures of the brain's pathways obtained by these new techniques led to the coining of the term **connectome** to indicate the "structural description of the network of elements and connections forming the human brain" (Sporns et al., 2005), or more simply, the "wiring diagram" of neurons in the brain (Baronchelli et al., 2013).

Determining the brain's wiring diagram is an important step in understanding how different areas of the brain communicate, because communication depends on structural connections. Interestingly, maps of structural connectivity of the brain have recently been likened to "fingerprints" that are different for every person, so it could be argued that the brain's wiring makes us who we are (Finn et al., 2015; Seung, 2012; Yeh et al., 2016). But to fully understand how the brain's structural network makes us who we are, or how it helps create cognition, it is necessary to determine how groups of neurons within the connectome form *functional* connections that are related to specific types of cognition.

Functional Connectivity

Picture the road network of a large city. On one set of roads, cars are streaming toward the shopping mall just outside the city, while other roads are funneling cars toward the city's business and financial district. One group of people is using roads to reach places to shop. Another group is using roads to get to work or to conduct business. Just as different parts of the city's road network are involved in achieving different goals, so different parts of the brain's neural network are involved in carrying out different cognitive or motor tasks.

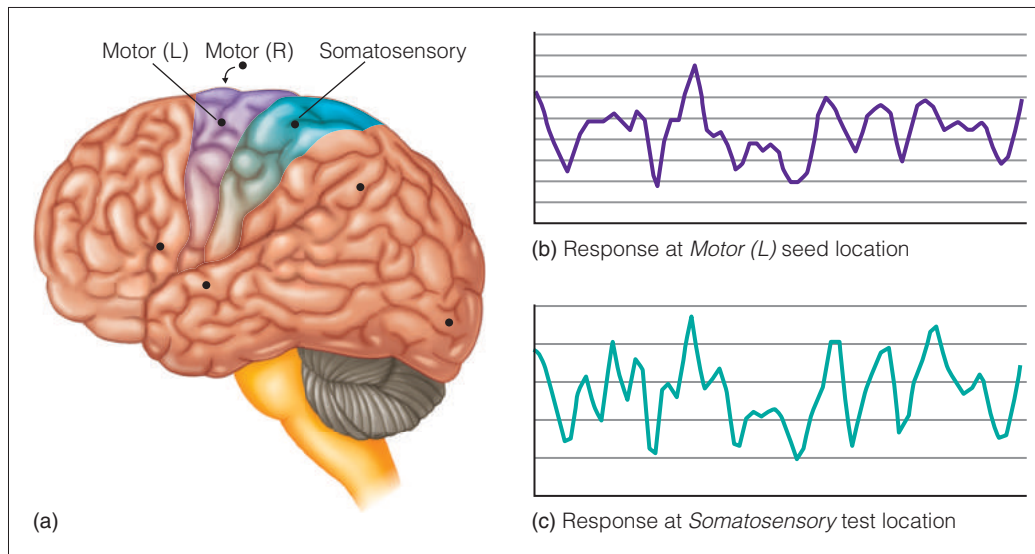
How is it possible to determine what parts of a neural network are involved in different functions? One way this question has been answered is by measuring **functional connectivity**, with functional connectivity being determined by the extent to which neural activity in two brain areas are correlated (Harmelech & Malach, 2013; Pessoa, 2014). If the responses of two brain areas are correlated with each other, this means that they are **functionally connected**.

One method of determining whether the responding of two areas is correlated is based on **resting-state fMRI**—the fMRI response measured while a person is at rest (that is, not performing a cognitive task). The procedure for measuring **resting-state functional connectivity** was introduced by Bharat Biswal and coworkers (1995).

METHOD Resting-State Functional Connectivity

Resting-state functional connectivity is measured as follows:

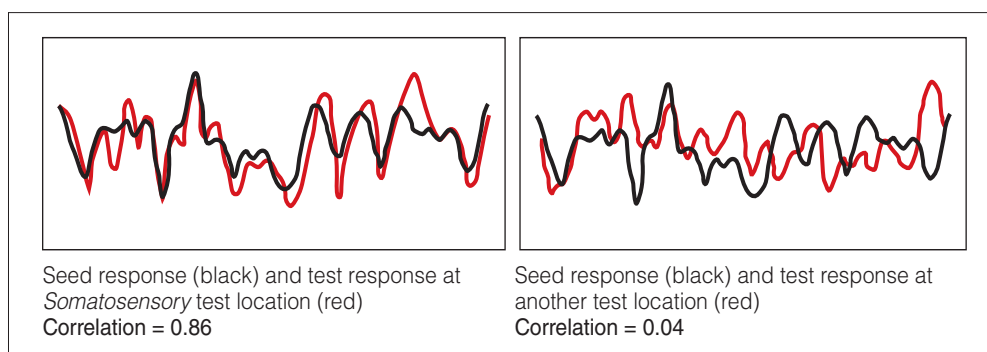
1. Use task-related fMRI to determine a brain location associated with carrying out a specific task. For example, movement of the finger causes an fMRI response at the location marked *Motor (L)* in **Figure 2.25a**. This location is called the **seed location**.



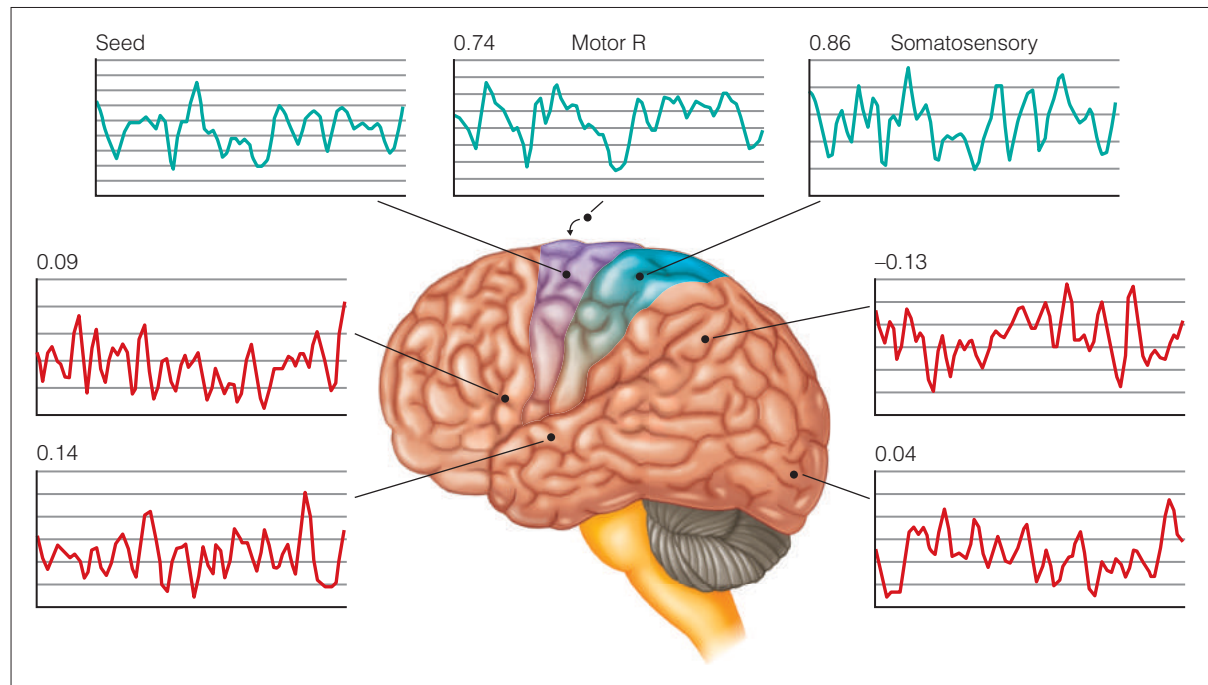
► **Figure 2.25** (a) Left hemisphere of the brain, showing the seed location *Motor (L)* in the left motor cortex, and a number of test locations, each indicated by a dot. Test location *Motor (R)* is in the right motor cortex on the other side of the brain from *Motor (L)*. Test location *Somatosensory* is in the somatosensory cortex, which is involved in perceiving touch. (b) Resting level fMRI response of the *Motor (L)* seed location. (c) Resting level fMRI response of the *Somatosensory* test location. The responses in (b) and (c) are 4 seconds long.

(Source: Responses courtesy of Ying-Hui Chou)

2. Measure the *resting-state fMRI* at the seed location. The resting-state fMRI of the seed location, shown in **Figure 2.25b**, is called a **time-series response** because it indicates how the response changes over time.
3. Measure the resting-state fMRI at another location, which is called the **test location**. The response of the test location *Somatosensory*, which is located in an area of the brain responsible for sensing touch, is shown in **Figure 2.25c**.
4. Calculate the correlation between the seed- and test-location responses. The correlation is calculated using a complex mathematical procedure that compares the seed and test responses at a large number of places along the horizontal time axis. **Figure 2.26a** shows the response at the *Somatosensory* test location superimposed on the seed response. The **correspondence between these responses results in a high correlation, which indicates high functional connectivity**. **Figure 2.26b** shows the seed response and the response at another test location. The **poor match between these two responses results in a low correlation, which indicates poor or no functional connectivity**.



► **Figure 2.26** Superimposed seed response (black) and test-location response (red). (a) Response of the *Somatosensory* test location, which is highly correlated with the seed response (correlation = 0.86). (b) Response of another test location, which is poorly correlated with the seed response (correlation = 0.04). (Source: Responses courtesy of Ying-Hui Chou)



► **Figure 2.27** Resting-state fMRI responses for the *Motor (L)* seed, test locations *Motor (R)*, *Somatosensory*, and five test locations in other parts of the brain. The numbers indicate correlations between the seed response and each test-location response. Responses *Motor (R)* and *Somatosensory* have been singled out because they have high correlations, which indicates high functional connectivity with the seed. The other locations have low correlations so are not functionally connected to the seed location.

(Responses courtesy of Ying-Hui Chou.)

Figure 2.27 shows time-series for the seed location and a number of test locations, and the correlations between the seed and test locations. The test locations, *Somatosensory* and *Motor (R)*, are highly correlated with the seed response and so have high functional connectivity with the seed location. This is evidence that these structures are part of a functional network. All of the other locations have low correlations so are not part of the network.

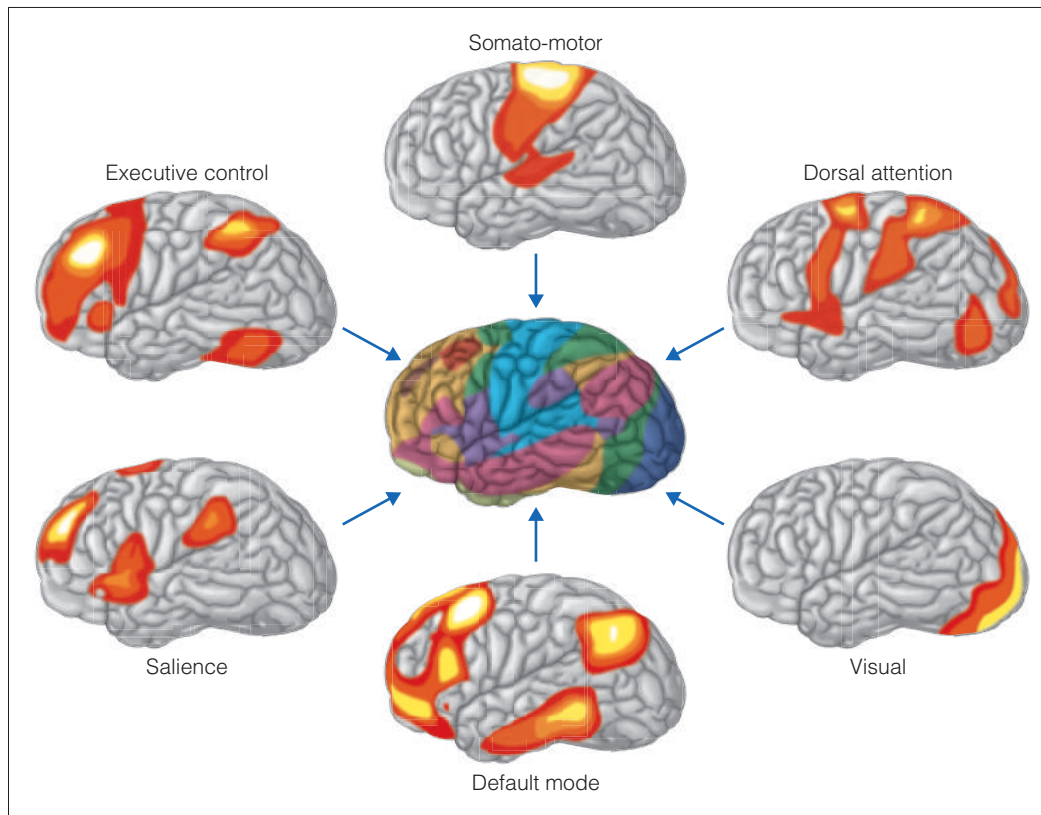
Resting-state functional fMRI connectivity has become one of the main methods for determining functional connectivity. **Figure 2.28** shows networks determined for a number of different functions using this procedure. **Table 2.1** summarizes the functions of these networks.

TABLE 2.1

Six Common Functional Networks Determined by Resting-State fMRI

Network	Function
Visual	Vision; visual perception
Somato-motor	Movement and touch
Dorsal Attention	Attention to visual stimuli and spatial locations
Executive Control	Higher-level cognitive tasks involved in working memory (see Chapter 5) and directing attention during tasks
Salience	Attending to survival-relevant events in the environment
Default mode	Mind wandering, and cognitive activity related to personal life-story, social functions, and monitoring internal emotional states

Sources: From Barch, 2013; Bressler & Menon, 2010; Raichle, 2011; Zabelina & Andrews-Hanna, 2016. Note that there are other networks as well, including networks involved in hearing, memory, and language.



► **Figure 2.28** Six major brain networks determined by the resting-state fMRI procedure. Note that all of these networks increase activity during a task and decrease activity when at rest, *except* the default mode network, which decreases activity during a task, and increases activity when there is no task. See Table 2.1 for brief descriptions of these networks. (Source: From Zabelina & Andrews-Hanna, 2016)

There are also other ways to determine functional connectivity. For example, functional connectivity can be determined by measuring the task-related fMRI at the seed and test locations and determining the correlations between the two responses, as we did for resting-state fMRI. **It is important to note that saying two areas are *functionally connected* does not necessarily mean that they directly *communicate* by neural pathways. For example, the response from two areas can be highly correlated because they are both receiving inputs from another area.** Functional connectivity and structural connectivity are not, therefore, the same thing, but they are related, so regions with high structural connectivity often show a high level of functional connectivity (Poldrack et al., 2015; van den Heuvel & Pol, 2010).

The pictures in Figure 2.28 show that the overall structural map of the brain is divided into smaller functional maps, so different cognitions activate different groups of neurons. But to really understand what is happening during cognition, we need to go beyond just identifying areas that serve different functions. We need to consider the *dynamics* of cognition.

The Dynamics of Cognition

To understand what we mean by the dynamics of cognition, let's return to our analogy between the structural map of the brain and a big-city street system. Imagine climbing into a helicopter and flying above the city so you can observe the patterns of traffic flow at various times of day. As you hover above the city, you notice how this flow changes when the street

system is serving different functions. During morning rush hour, when its function is to get people to work, there is heavy flow from the suburbs toward the city on the major highways. Evening rush hour reverses the flow on the major highways, as people head for home, and the flow may also increase on suburban streets a little later. During the day, traffic flow may be higher around shopping areas; and before and after special events, like a weekend football game, flow will be high on roads leading to and from the stadium.

The point of this example is that just as traffic flow in the city changes depending on conditions, the flow of activity within and across the functional networks in the brain also changes, depending on conditions. For example, let's consider what is happening when a person looks at a cup of coffee on a table. Looking at the cup causes activity in the visual functional network, as the person perceives the various qualities of the cup. Meanwhile, the attention network may also be activated, as the person focuses attention on the cup, and then the motor network becomes activated as the person reaches to pick up the cup, grasps it, and lifts it to drink. So even a simple everyday experience like looking at and picking up a cup of coffee involves rapid switching and sharing of information between a number of different functional networks (van den Heuvel & Pol, 2010).

In addition to this rapid switching between networks, changes in connectivity can also occur more slowly. For example, functional connectivity changes in memory networks from morning to evening as memories are accumulated during the day and are then strengthened at night (Shannon et al., 2013). Connectivity changes have also been reported to occur in response to eating food or drinking coffee, with some networks being strengthened and others weakened when a person who had fasted for a day resumed eating and drinking (Poldrack et al., 2015; also see McMenamin et al., 2014 for the effect of anxiety on networks). Functional networks are not, therefore, simply static diagrams but involve constantly changing activity within and across networks (Mattar et al., 2015).

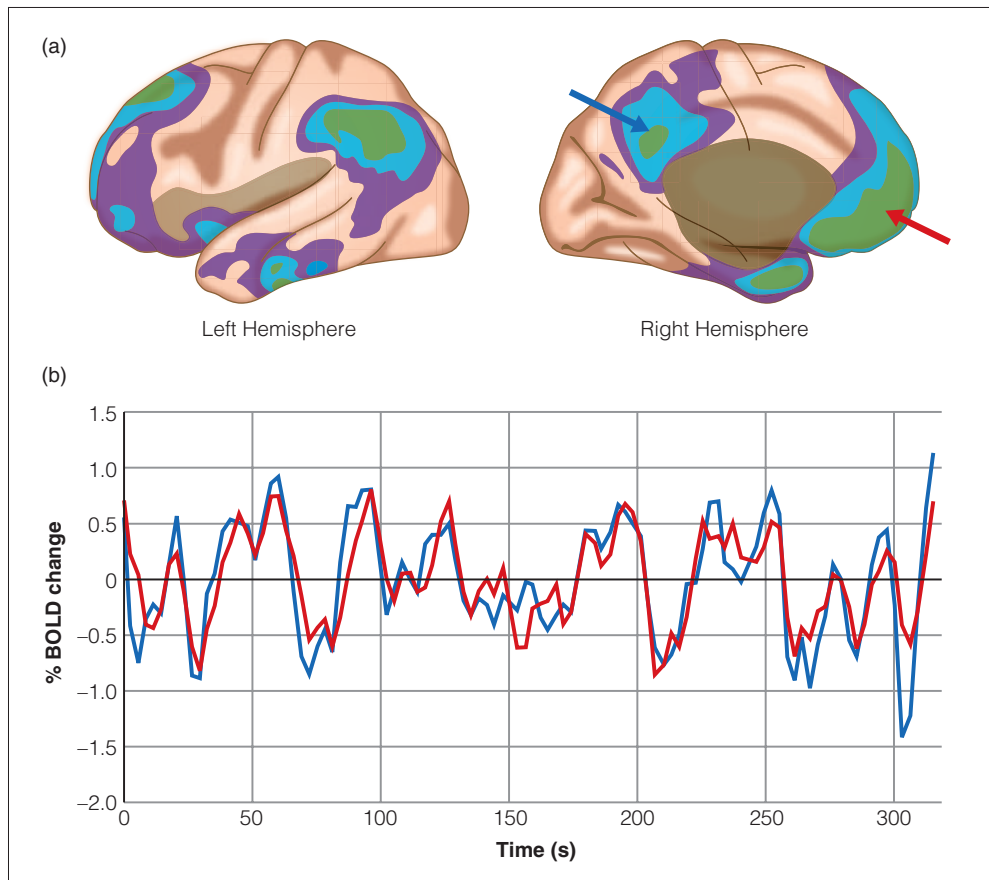
Many of the ideas about networks and connectivity that we have been describing are based on recent research that logically follows from the idea of distributed representation. After all, if functions are represented by structures in many different areas of the brain, it makes sense that they would be able to communicate with each other. But in the process of discovering ideas about how networks work, one finding, made in just the last two decades, was extremely unexpected. A network was discovered that responded not when people were engaged in tasks—but when they weren't! This network is called the *default mode network*.

The Default Mode Network

The **default mode network (DMN)**, which is shown at the bottom of Figure 2.28, is a network of structures that respond when a person is not involved in specific tasks. The story behind the discovery of this network begins with a paper by Gordon Shulman and coworkers (1997), who noted a few earlier fMRI studies in which presentation of a task caused a *decrease* in activity in some areas of the brain, and stopping the task caused an increase in activity in the same areas. This was different than the usual result, in which presentation of a task is associated with an increase in activity and stopping the task is associated with a decrease in activity.

Following up on this observation, Marcus Raichle and coworkers (2001), in a paper titled “A Default Mode of Brain Function,” proposed that the areas that decrease activity during tasks represent a “default mode” of brain function—that is, a mode of brain function that occurs when it is at rest.

To make things even more interesting, research using the resting-state functional connectivity method indicated that areas in the frontal and parietal lobes that decrease activity during tasks (Figure 2.29a) have correlated resting state activity (Figure 2.29b)



► **Figure 2.29** (a) Brain areas that decrease their activity during task performance. (b) Resting-state activity at two points in the right hemisphere, indicated by the arrows above. The fact that the resting state activity is correlated indicates that these areas are functionally connected. All of these areas, taken together, are called the *default mode network*. (Source: From Raichle, 2015)

(Greicius et al., 2003). These areas are, therefore, part of a functional network, which is identified as the default mode network (DMN) in Figure 2.28.

There has been a great deal of speculation and research about the purpose of the DMN. **One interesting observation is that when the DMN is active, people's minds tend to wander (Killingsworth & Gilbert, 2010; Smallwood & Schooler, 2015).** This is probably something you have experienced. One minute you are driving along the highway, paying close attention to your driving, but then, without even realizing it, you find your mind wandering to thoughts about what you are going to do later or how to deal with some on-going concern. What happened? Your brain switched from task-related networks involved in driving to your default mode network. As you might expect, mind wandering during driving, as well as during other tasks, isn't necessarily a good thing for performance. This idea is confirmed by a large body of research that shows that mind wandering decreases performance on tasks that require focused attention (Lerner, 2015; Moneyham & Schooler, 2013; Smallwood, 2011).

But the DMN must have a purpose other than creating mind wandering that distracts you from your work! After all, it is one of the brain's largest networks and accounts for a large portion of the brain's activity when it is at rest. In some of the chapters that follow, we will consider evidence that shows that the DMN is involved in processes ranging from attention to memory to creativity.

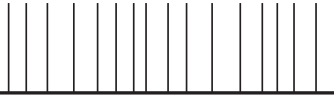
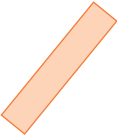
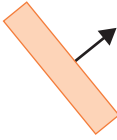






► SOMETHING TO CONSIDER

Technology Determines the Questions We Can Ask

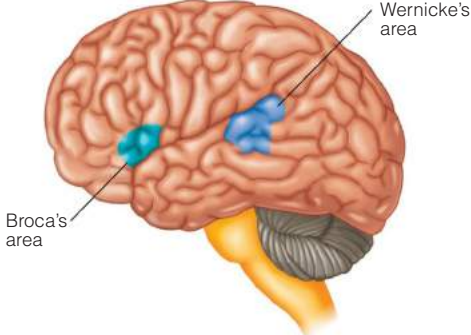
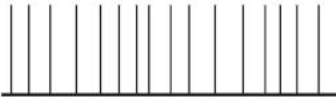
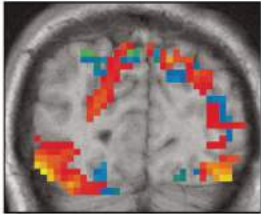
We've covered a lot of ground in this chapter, all the way from early neuropsychological research on how brain damage affects language, to how single neurons respond to visual stimuli, to how multidimensional cognitions are represented by dynamic activity in interconnected neural networks.

When we consider all of the methods we have described, we can see that the questions that researchers have asked depended on the available technology. Consider, for example, *The Representation Question*: “How are cognitions represented by neural firing?” (Figure 2.30). The first step toward being able to answer that question was the introduction of the technique for recording from single neurons in 1928. But it wasn't until later, beginning in the 1950s, when more-advanced electrodes and amplifiers made it possible to record from single neurons in the brain. As soon as that became possible, researchers were able to move from asking “how do neurons respond to a flash of light?” to “how do neurons respond to complex shapes?”

The quest to determine how neurons respond to different kinds of visual stimuli had an additional effect. It led researchers to begin recording from neurons in areas of the brain outside of the visual cortex. These forays to other areas of the brain, such as the temporal cortex, revealed that as much as half of the brain is activated by visual stimuli (Van Essen, 2004). Later research showed that other cognitive functions, such as hearing, pain, memory, and language, activate numerous areas of the brain.

The Representation Question	
How are cognitions represented by neural firing?	
Method: Single neuron recording	
1928: Adrian's first single neuron recording (but not the brain)	
1960's: Hubel and Wiesel's feature detectors in the brain	<div style="display: flex; justify-content: space-around; align-items: center;"> <div style="text-align: center;">  <p>Oriented bar</p> </div> <div style="text-align: center;">  <p>Oriented moving bar</p> </div> <div style="text-align: center;">  <p>Short moving bar</p> </div> </div>
1970's: Neurons that respond to complex stimuli	<div style="display: flex; justify-content: center; align-items: center;">      </div> <div style="display: flex; justify-content: center; gap: 20px; margin-top: 5px;"> 3 4 4 5 6 </div>

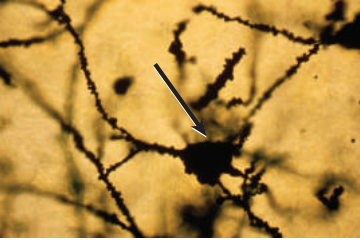
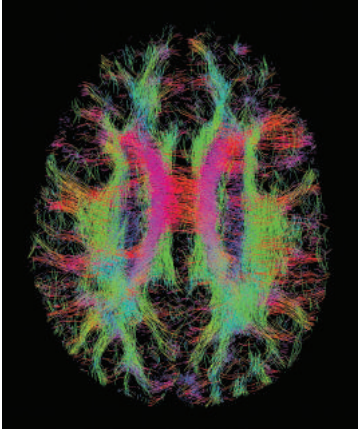
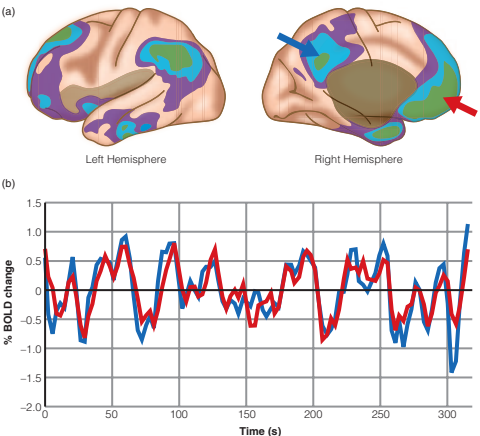
► **Figure 2.30** Technology involved in research studying how cognitions are represented by neural firing. Advances in technology made it possible to record from the brain and study how neurons respond to more complex stimuli.

The Organization Question	
How are cognitions localized in different places in the brain?	
Method: Neuropsychology: Study of behavior of people with brain damage	
1860s: Broca and Wernicke, followed later by modern research	
Method: Single neuron recording; record from different areas of the brain	
1960s on: Hubel and Wiesel's feature detectors in the brain 1970s on: Neurons that respond to complex stimuli	
Method: Brain imaging	
1976: First PET scan 1990: First fMRI	 <p style="font-size: small; text-align: right;">Source: From Ishai et al., 2000</p>

► **Figure 2.31** Technology involved in research studying how cognitions are localized in different areas of the brain. Three different methods have been used: Neuropsychology (1860 on); Single neuron recording (1960s on); and Brain imaging (beginning in 1976 with PET, picking up steam in 1990, with the introduction of fMRI).

The discovery that stimuli can activate a large area of the brain brings us to *The Organization Question*: “How are cognitions localized in different areas of the brain?” (Figure 2.31). This question was first studied in humans in the 1800s by Broca and Wernicke’s research on brain-damaged patients and then in animals in the 1960s and 70s using single neuron recording. Although a tremendous amount was learned about brain organization using these techniques, research on human brain organization took off when brain-scanning technology was introduced—first PET scans in 1976 and then fMRI in 1990, which made it possible to map patterns of brain activity in humans.

But just determining which brain areas were activated wasn’t enough for researchers. They wanted to go beyond determining static maps to study dynamic communication, and so they posed *The Communication Question*: “How are different areas of the brain connected and how do they communicate?” (Figure 2.32). The idea that neurons form circuits

The Connectivity Question	
How are different areas connected and how do they communicate?	
Structural Connectivity	
Method: Classical neuroanatomy	
1870s on	 <p style="font-size: small; text-align: right;">Clouds Hill Imaging Ltd./Corbis Documentary/Getty Images</p>
Method: Track weighted imaging and others	
1990s on	 <p style="font-size: small; text-align: right;">Courtesy of Fernando Calamante and Elsevier</p>
Functional Connectivity	
Method: Resting state fMRI	
1995: First resting state fMRI	 <p style="font-size: small;">(a) Left Hemisphere Right Hemisphere</p> <p style="font-size: small;">(b) % BOLD change Time (s)</p>

► **Figure 2.32** Technology involved in research studying how different areas of the brain are connected and how they communicate. Structural connectivity was studied in the 1800s using neuroanatomical techniques and beginning in the 1990s using brain imaging. The study of functional connectivity took off with the introduction of the resting-state fMRI method in 1995.

dates back to early anatomy experiments from the 1800s, which, like the single neuron recordings, were carried out on animals. But the introduction of brain imaging and other technologies made it possible to begin determining structural connectivity (the “roadmap” of the brain) and functional connectivity (the “traffic pattern” of the brain).

The events we have been describing show that technology has determined not only what can be learned about the functioning of the brain but also the types of behaviors that can be studied. Early research involved simple behaviors—the ability to perceive a flash of light, an oriented line, or a geometrical shape. Even later, when researchers began presenting more complex objects like faces, they were usually presented as briefly flashed pictures. But present-day research involves more naturalistic stimuli, such as events depicted in films. And perhaps more important, although early research focused largely on visual stimuli, current research has expanded to include cognitive behaviors ranging from remembering the past and imagining the future to understanding sentences and making decisions.

But before we get too carried away by the wonders of technology, let’s not lose sight of the fact that although it may be nice to know how neurons work, where brain structures are located, or how neurons communicate in networks, **psychologists are not really interested in studying physiology for its own sake.** They are interested in determining the relationship between physiological mechanisms and experiences, thoughts, and actions.

The approach in this book is, therefore, based on the idea that the best way to explain cognition is by conducting both behavioral and physiological experiments. As you read this book, you will encounter many examples of situations in which the results of behavioral and physiological experiments have been used together to provide a richer understanding of the mind than would be provided by either alone.

TEST YOURSELF 2.2

1. What is localization of function? Describe how localization has been demonstrated by neuropsychology and recording from neurons. Be sure you understand the principle of double dissociations.
2. Describe the basic principles behind functional magnetic resonance imaging.
3. Describe brain-imaging evidence for localization of function. Describe experiments that involved looking at still pictures and that involved looking at movies. What does each type of experiment tell us about localization of function?
4. What is distributed representation? How is distributed representation related to the multidimensional nature of experience? How is distributed processing illustrated by how the brain responds to looking at faces, remembering and language?
5. What is a neural network?
6. What is structural connectivity? How is it measured?
7. What is functional connectivity? How is it measured and what are some networks that have been determined using this technique?
8. What does it mean to say that the operation of brain networks is dynamic?
9. What is the default mode network? How is it different than other networks?
10. Describe the connection between advances in technology and research on the physiology of cognition.

CHAPTER SUMMARY

1. Cognitive neuroscience is the study of the physiological basis of cognition. Taking a levels-of-analysis approach to the study of the mind involves research at both behavioral and physiological levels.
2. Ramon y Cajal's research resulted in the abandonment of the neural net theory in favor of the neuron doctrine, which states that individual cells called *neurons* transmit signals in the nervous system.
3. Signals can be recorded from neurons using microelectrodes. Edgar Adrian, who recorded the first signals from single neurons, determined that action potentials remain the same size as they travel down an axon and that increasing stimulus intensity increases the rate of nerve firing.
4. The principle of neural representation states that everything that a person experiences is based not on direct contact with stimuli, but on representations in the person's nervous system.
5. Representation by neurons can be explained by considering feature detectors, neurons that respond to complex stimuli, and how neurons are involved in specificity coding, population coding, and sparse coding.
6. The idea of localization of function in perception is supported by the existence of a separate primary receiving area for each sense, by the effects of brain damage on perception (for example, prosopagnosia), by recording from single neurons, and from the results of brain-imaging experiments.
7. Brain imaging measures brain activation by measuring blood flow in the brain. Functional magnetic resonance imaging (fMRI) is widely used to determine brain activation during cognitive functioning. Brain-imaging experiments have measured the response to still pictures to identify areas in the human brain that respond best to faces, places, and bodies, and the response to movies to create a brain map indicating the kinds of stimuli that activate different areas of the brain.
8. The idea of distributed processing is that specific functions are processed by many different areas in the brain. One reason for the activation of many areas is the multidimensional nature of experience. This principle is illustrated by the multidimensional nature of seeing a face, remembering, and producing and understanding language.
9. Neural networks are groups of neurons or structures that are connected structurally and also that are functionally related.
10. Structural connectivity defines the neural highway system of the brain. It has been measured using track weighted imaging.
11. Functional connectivity occurs when different areas have temporally correlated responses. Measuring resting-level fMRI has emerged as one of the ways to measure functional connectivity, but functional connectivity can also be measured by task-related fMRI.
12. A number of different functional networks, such as visual, auditory, salience, executive function, and motor networks, have been determined using resting-level fMRI.
13. A full description of networks needs to include the dynamic aspects of network activity.
14. The default mode network is different than other networks because its activity decreases when a person is engaged in a task, but then increases when the brain is at rest. The function of the DMN is still being researched, but it has been suggested that it may play important roles in a number of cognitive processes, which we will discuss later in the book.
15. Progress in understanding the physiology of cognition has depended on advances in technology. This is demonstrated by considering the connection between technology and answering three basic questions: The Representation Question, The Organization Question, and The Communication Question.

THINK ABOUT IT

1. Some cognitive psychologists have called the brain the mind's computer. What are computers good at that the brain is not? How do you think the brain and computers compare in terms of complexity? What advantage does the brain have over a computer?
2. People generally feel that they are experiencing their environment directly, especially when it comes to sensory experiences such as seeing, hearing, or feeling the texture of a surface. However, our knowledge of how the nervous system operates indicates that this is not the case. Why would a physiologist say that all of our experiences are indirect?
3. When brain activity is being measured in an fMRI scanner, the person's head is surrounded by an array of magnets and