This excerpt from

Cognitive Science: An Introduction - 2nd Edition.
Neil A. Stillings, Steven E. Weisler, Christopher H. Chase,
Mark H. Feinstein, Jay L. Garfield and Edwina L. Rissland.
© 1995 The MIT Press.

is provided in screen-viewable form for personal use only by members of MIT CogNet.

Unauthorized use or dissemination of this information is expressly forbidden.

If you have any questions about this material, please contact cognetadmin@cognet.mit.edu.

Chapter 7
Neuroscience: Brain and Cognition

7.1 Introduction to the Study of the Nervous System

In the previous chapters we have taken a computational approach to cognition, analyzing information processes as computations over abstractly defined representations. Although human and animal cognition is physically realized in the nervous system, the computational level of analysis has allowed us to study cognition without paying much attention to its physical implementation. In this chapter we introduce *neuroscience*, the field that encompasses the levels of analysis that are required for the study of physical processes and structures in the nervous system. There are several reasons to extend our study of cognition to include its biological underpinnings.

First, the question of how nervous systems achieve the kinds of complex information processing that we have been studying is intrinsically fascinating. Questions about how to complete the chain of scientific understanding from mind to brain, from thought to neuron to molecule and perhaps on to subatomic particle, are among the most interesting in science.

Second, research on the nervous system can help test some of the theories that have been developed in cognitive psychology and linguistics. We saw in chapter 2, for example, how the discovery that visual areas of the brain are differentially active during imagery is an important piece of evidence for the existence of a spatially organized buffer for visual imagery.

Third, as we saw in our study of connectionism, the potential importance of neuroscience to cognitive science goes beyond the likelihood that neuroscience is a source of evidence for autonomously developed computational theories. The information processes that a system is capable of carrying out efficiently are strongly constrained by its computational architecture. The capabilities of an architecture arise directly out of its physical structure. Thus, knowledge about fundamental principles of structure and process in the nervous system should be able to contribute to the initial construction of a theory of cognitive architecture. Ideally, a theory of the architecture would be the joint product of findings at the computational and biological levels of analysis. Connectionist researchers have reached for this ideal by trying to build principles of neural computation into their conceptions of cognitive architecture. There is considerable controversy about whether the principles that have been suggested are the right ones and even whether the whole enterprise is premature, given our primitive understanding of neural computation, but the vision cannot be faulted. In order to convey a sense of what is known in neuroscience and how it relates to analysis at the cognitive level, this chapter covers quite a bit of territory. Some of what is presented falls into the area that has come to be called *cognitive neuroscience*, the active intersection between the two fields. Other material is currently rather remote from the central

concerns of cognitive science, but could become more important as research advances. The cognitive scientist should be conversant with neuroscience, ready to establish new connections when they become possible.

A fourth reason for linking the study of cognitive science and neuroscience is that the understanding of the relationship between the physical structure of the nervous system and its information-processing capacities has many potential practical applications. Increased understanding of the biological substrates of cognition will contribute to the development of better physical and behavioral treatments for damage to the nervous system caused by accident or disease. More generally, such knowledge will aid in the design of learning environments that are tuned to the physically determined strengths of human information processing and will suggest new ways in which computer systems can be used to compensate for its weaknesses. Researchers have also begun to envision new types of computers with highly parallel hardware that exploits some of the design principles of the nervous system (Hecht-Nielsen 1990).

7.2 *Organization of the Central Nervous System*

Introduction: Levels of Description

The nervous system has been studied at many different levels. Neurobiology gives a very different picture of what the brain is doing than psycholinguistics or neuropsychology. As Sejnowski and Churchland (1989) have pointed out, these divisions of study are somewhat arbitrary and serve the convenience of scientists and the techniques that are used in their research. They describe seven different levels (see figure 7.1), organized along a spatial scale. At the top of the scale are neural systems organized around general functional characteristics, such as the speech articulation system. At the bottom of the scale are neurotransmitters, the chemicals that carry signals between neurons, the cells in the brain that process information. At present our

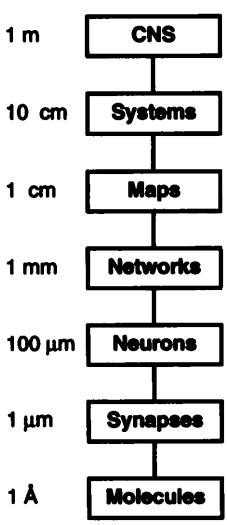

Figure 7.1
Scale of size and unit of function in the nervous system

knowledge of neural processes at the molecular or synaptic level is much more detailed than at the neural network and map levels. Future research is likely to restructure these descriptive levels as we learn more about the brain's own states of internal representation.

Basic Neuroanatomy

The human nervous system is divided into central and peripheral systems. The *central nervous system* is made up of the brain and spinal cord and can be thought of as the control center for interpreting sensory input and directing our thoughts and actions. The *peripheral nervous system* functions to carry information from the body and the outside world to the central system and back. For understanding cognitive functions, we are most interested in a review of the central nervous system (CNS).

Neuroanatomy refers to the general structural organization of the brain and spinal cord, including its major physical and functional divisions. Neuroanatomy is difficult to learn because new brain structures developed on top of older ones during the course of evolution, producing a highly complicated organ with many regions tightly packed together and extremely complex interconnected pathways among them. Figure 7.2 illustrates the major regions of the CNS and some of its principal structures. The lowest functional division is the spinal cord, which connects the brain to the body's organs and muscles. The middle division is the *brain stem*, which is made up of the diencephalon (between brain), the midbrain or mesencephalon, and the hindbrain or rhombencephalon. The highest division is the *forebrain*, which includes the cerebral cortex, basal ganglia, olfactory bulbs, and limbic system. There are complementary left and right structures at every division in the brain, with the exception of the pineal body.

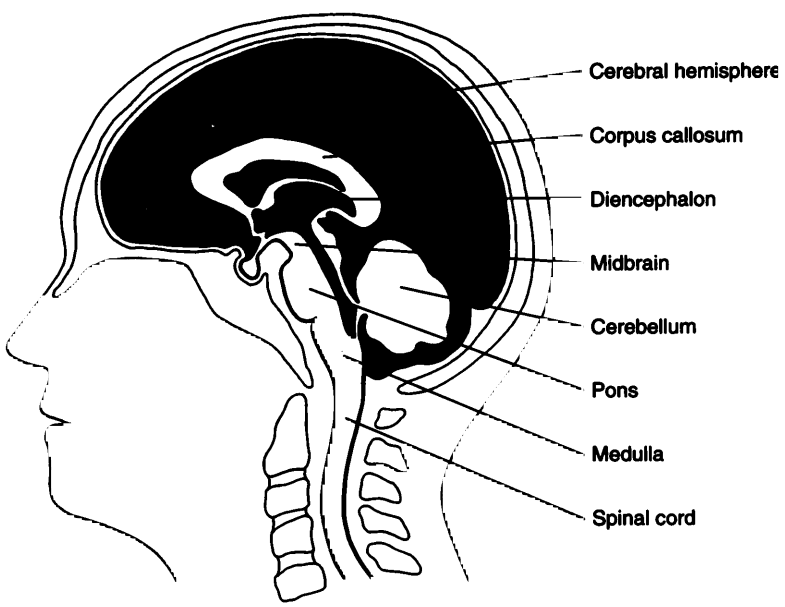

Figure 7.2
A midsagittal diagram of the principal structures of the brain. (From Kandel 1991a.)

The diencephalon consists of three thalamic (inner chamber) structures and the pituitary gland. The epithalamus is made up of the pineal body and other cell groups, or nuclei, called the habenula. The hypothalamus is composed of over twenty small nuclei. They are involved in nearly all aspects of behavior, including regulatory functions of sleep, appetite, reproductive cycles, and body temperature. These functions are affected through the release of chemicals called hormones into the blood stream that are then transported to other tissues in the body and brain. The *thalamus*, made up of a cluster of smaller nuclei, is located dorsally (on top) and anterior (in front of) to the rest of the midbrain. The thalamus functions as part of the sensory relay system that receives input from all sensory organs except the nose and relays these signals to the cortex. For example, the *lateral geniculate nucleus* (LGN) receives input from the eyes, and the *medial geniculate nucleus* (MGN) receives auditory projections.

The midbrain is organized into two sections that are physically divided by the cerebral aqueduct. This narrow cavity is part of the ventricle system in the brain, which is filled with cerebrospinal fluid, cushioning the brain from shock and possibly playing a role in filtering metabolic wastes. The area lying above the aqueduct is called the tectum and primarily consists of the superior and inferior colliculi, nuclei that receive projections from visual and auditory organs, respectively. Below the aqueduct is the tegmentum, made up of sensory and motor fibers passing between the forebrain and the peripheral nervous system, and a number of motor nuclei, such as the substantia nigra (which when damaged produces the symptoms of Parkinson's disease).

The hindbrain is made up of the pons, medulla oblongata, and cerebellum. In the pons are a variety of sensory and motor nuclei that govern vestibular (balance and postural orientation) and motor functions. The medulla oblongata consists primarily of fiber tracts that pass information between the spinal cord and the cortex. In addition, it contains a complex mixture of fibers, called the reticular formation, that travel up and down the brain stem from the diencephalon through the hindbrain.

Overlying the brain stem is the *cerebellum*. Its cellular organization is remarkably uniform throughout, as compared, for example, to the cortex, which has regionally organized characteristics. The cerebellum has connections to structures throughout the rest of the brain. Although it was once thought to be specialized primarily for sensory-motor coordination, recent research suggests roles in learning (McCormick and Thompson 1984) and possibly other higher cognitive functions (Courchesne et al. 1988).

The forebrain is divided into four regions: the cerebral cortex, the limbic system, olfactory bulbs, and basal ganglia. Sometimes the thalamus is placed with the forebrain group. The limbic (border) system is made up of several structures. The major nuclei include the hippocampus (seahorse), amygdala, septum (partition), mammillary bodies, fornix, and cingulate (girdle) gyrus, which together surround and sheathe the brain stem. These structures form connections with parts of the hypothalamus, thalamus, and cortex through the cingulate gyrus. Because of their interconnections with the olfactory bulbs they were once called the rhinencephalon and thought to analyze olfactory information. More recently the hippocampus has received intense study for its role in the formation of memory (Squire 1987) and in spatial orientation (O'Keefe and Nadel 1978).

The basal ganglia are formed from several large nuclei that surround the thalamus and lie in the medial (middle) region below the cortex. These structures, including the putamen (shell), the globus pallidus, and the caudate, have extensive connections with

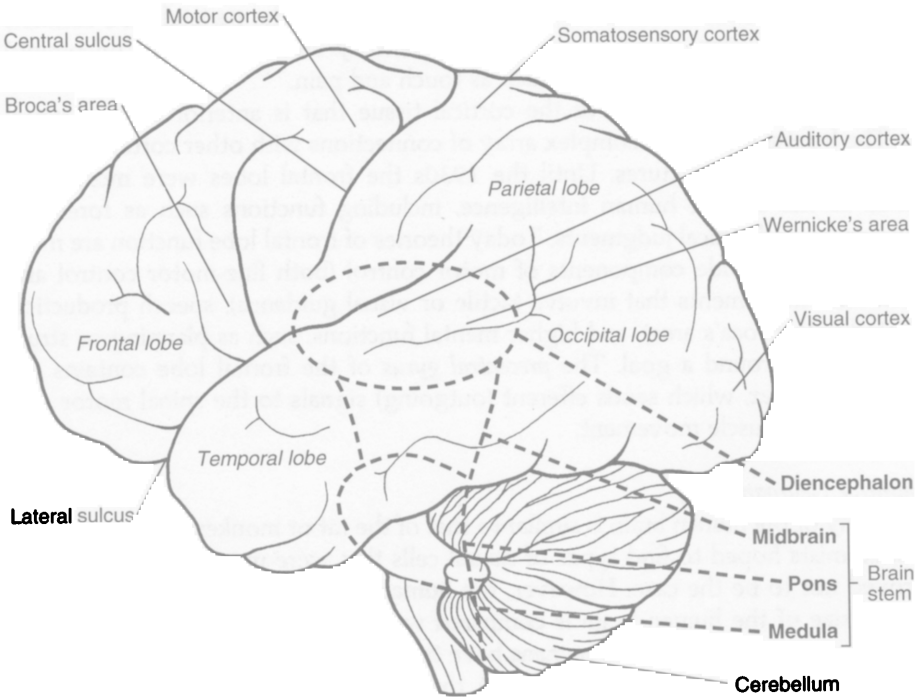

Figure 7.3
The main parts of the central nervous system and cerebral cortex. (From Kandel 1991a.)

the cortex and thalamus as well as portions of the midbrain, such as the red nucleus and substantia nigra. They play a significant role in the initiation and control of movement.

Cerebral Cortex Approximately 1.5–3 mm thick, the *cerebral cortex* consists of four to six layers of cells that are spread out over a large surface area of 100,000–200,000 mm^2 (Cherniak 1990). As the cortex increased in size during the course of evolution, folds and wrinkles appeared, presumably to allow more cortical tissue to fit into the skull without expanding the size of the cranium. This folding process produced characteristic ridges (called gyri or lobules) and clefts (called fissures or sulci if shallow).

These ridges and clefts are not identical in any two individuals, but they are relatively uniform and can be divided into four major regions (see figure 7.3). The *occipital lobes* are located in the posterior (back) part of the cortex. They serve as the primary sensory system for vision. The *temporal lobe*, located lateral (to the side) and ventral (below) to the prominent lateral or *Sylvian fissure*, serves at least three functions: primary and secondary processing of auditory information, long-term storage of sensory information (such as the identity of visual objects) (Mishkin and Appenzeller 1987), and processing of emotional qualities associated with sensation. The *parietal lobe* is strategically located posterior to the *central sulcus* and anterior to the occipital lobe. Its multiple connections with occipital, frontal, and temporal lobes suggest a functional role in the cross-modal integration of sensory information. Recent research also has suggested a functional role in locating visual objects and directing visual attention

(Andersen, Siegel, and Essick 1987). The *postcentral gyrus* of the parietal lobe contains the *somatosensory cortex*, which is the primary sensory area for afferent (incoming) signals from the surface of the body, such as touch and pain.

The *frontal lobe* is composed of the cortical tissue that is anterior to the central sulcus. The frontal lobe has a complex array of connections with other cortical regions and with subcortical structures. Until the 1930s the frontal lobes were most often described as the seat of human intelligence, including functions such as foresight, self-awareness, and ethical judgments. Today theories of frontal lobe function are more complicated and include components of motor control (both fine-motor control and complex limb movements that involve tactile or visual guidance), speech production (associated with Broca's area), and higher mental functions, such as planning or structuring activities around a goal. The *precentral gyrus* of the frontal lobe contains the *primary motor cortex*, which sends efferent (outgoing) signals to the spinal motor systems, producing muscle movement.

Comparative Neuroanatomy

The anatomy of the human brain is similar to that of the rat or monkey. Although early neuroanatomists hoped to find types of neural cells that were unique to humans, this has proved not to be the case. However, the human brain is different in other ways. When the size of the human brain is compared with the brain size of other animals, taking body weight into account, humans have the largest brains of all mammals—six times as large as a cat's brain and almost three times as large as a chimpanzee's (dolphins rank a close second to humans) (Stephen, Bauchot, and Andy 1970). Humans fare even better when the relative size of the cortex is used for comparison.

Compared with other primates, humans are born with a brain that is small relative to its adult size. Macaque monkeys are born with brains that are 60 percent of their adult weight. The proportions for chimpanzees and humans are 46 percent and 25 percent, respectively. These relationships appear to be due to the rate of brain growth that occurs after birth. During fetal development the brain grows at the same rate in macaque, chimpanzee, and human. At birth the rates of brain growth slow down markedly for macaque and chimpanzee, but human brains continue to grow at the rapid fetal rate for about two more years. Apparently, the increase in the size of the human brain was achieved evolutionarily by prolonging development after birth.

The significance of the facts about brain size is in some dispute. Differences in size could be the result of more cells, more cell processes (such as the size of dendritic branches or the number of axonal connections between neurons), or increased cell density. Using one method to estimate the density of cortical cells, Changeux (1985) reported that cell density and the relative frequency of different types of cells for a given cortical region are fairly constant across a number of species from mouse to human. He suggested that the human advantage is in the sheer number of cells, estimating humans to have from three to four times the number of cortical cells as other primates. Using other methods, Passingham (1982) calculated that with larger brain volume, human cortical cells are spaced farther apart than the cells in the brains of other primates. He suggested that in certain cortical areas the human brain has more connections per cell rather than more cells. The increased volume is not found in primary sensory regions of the cortex, such as the occipital lobe or parts of the temporal lobe, suggesting that the main evolutionary pressure on brain development

in humans was to enlarge cortical areas involved in the cross-modal integration of information, such as the parietal and frontal cortex.

The functional organization of the human brain also may be different from that of the brains of other mammals. For example, Passingham (1982) has suggested that because certain behavioral functions appear to be localized in one cortical hemisphere, such as language in the left hemisphere, humans may have a brain that can process information more efficiently. We will return to this topic later in the chapter.

Cellular Systems

Neurons The cells of the brain are generally divided into two categories: the nerve cells or *neurons* and others called *glial* cells. Neurons are intensively studied because of their unique computational properties. Although their biological processes are much like those of other cells, neurons have the special ability to receive and transmit information to other neurons, muscles, and glands, sometimes over great distances. Most neurons send signals to many other neurons and also receive signals from many other neurons.

Glial cells, which are much more numerous than neurons, are not well understood, but they are known to serve several functions. They are involved in the removal of unnecessary or excess substances. Glial cells often absorb excess neurotransmitter chemicals at synapses (a process described below) and have been observed to multiply and remove cellular debris at sites of brain damage. In vertebrates (animals with spinal cords) glial cells provide two special functions. They establish the blood-brain barrier to filter the blood supply to the brain and also form myelin on the axons of some neurons (both functions are described below in more detail).

The brain is the most metabolically active organ in the body, accounting for 15 to 20 percent of the body's oxygen utilization but only 2 percent of total body weight. Neurons are always active and require a constant supply of energy in the form of blood glucose to remain alive. As a neuron's level of activity increases, its need for energy rises as well and glucose is taken up more rapidly. Special imaging techniques, such as *positron emission tomography* (PET), have been developed to provide pictures of the brain by labeling glucose with a radioactive isotope. The labeled glucose is intravenously injected and differentially absorbed by those regions of the brain that are metabolically most active during the time period when the isotope is still active. Once the isotope has been absorbed into the cells, it remains there, and the subject's brain can be imaged without harmful effects. PET scans are one of the few procedures available that can provide functional brain images in humans. This technique is particularly useful for localizing behavioral functions to specific brain locations. It requires finding a task that involves areas of the brain differentially and keeping a subject at the task for two minutes without distraction or interruption while the image is being recorded.

If the supply of oxygenated blood to the brain is interrupted, through a heart attack or stroke, for example, the loss of oxygen disrupts cell metabolism and neurons begin to die within minutes. Because the genetic material in the soma of adult neurons cannot initiate the process for cell duplication (unlike glial cells), neurons cannot reproduce to replace lost cells. Thus, any brain damage in adults is permanent. Neurons are vulnerable in other ways as well. Their functioning can be disrupted by toxic substances from the environment or even by naturally occurring chemicals in the bloodstream such as

amino acids. For protection, a blood-brain barrier has evolved that acts to selectively filter the brain's blood supply. Glial cells, which have fatty membranes, tightly sheathe most of the blood vessels in the brain. Small molecules, such as oxygen, can easily pass through this membrane tissue, but larger ones, such as glucose, require special transport mechanisms to get across the barrier. For a drug to directly affect brain function, its molecular structure must be small enough to fit through the glial membrane or be soluble in the fatty tissue and so transportable across the membrane. In a few locations in the brain the blood vessels are not wrapped in glial tissue, which allows these brain regions to monitor the composition of the blood for substances such as hormones.

Although neurons vary considerably in shape and size (see figure 7.4), they generally have three major parts: the *cell body* or *soma*, the *dendrites* (Greek for tree), and the *axon*. The soma, whose diameter is usually smaller than 50 micrometers (one-millionth of a meter), contains the cell's nucleus and many other small structures (called organelles) that carry out processes necessary for the cell's health. Radiating out from the cell body are the dendrites and the axon. Each type of neuron has many dendritic fibers, which form a characteristic shape that distinguishes it from other types. For example, the dendrites of the Purkinje cells in the cerebellum have a complex coral

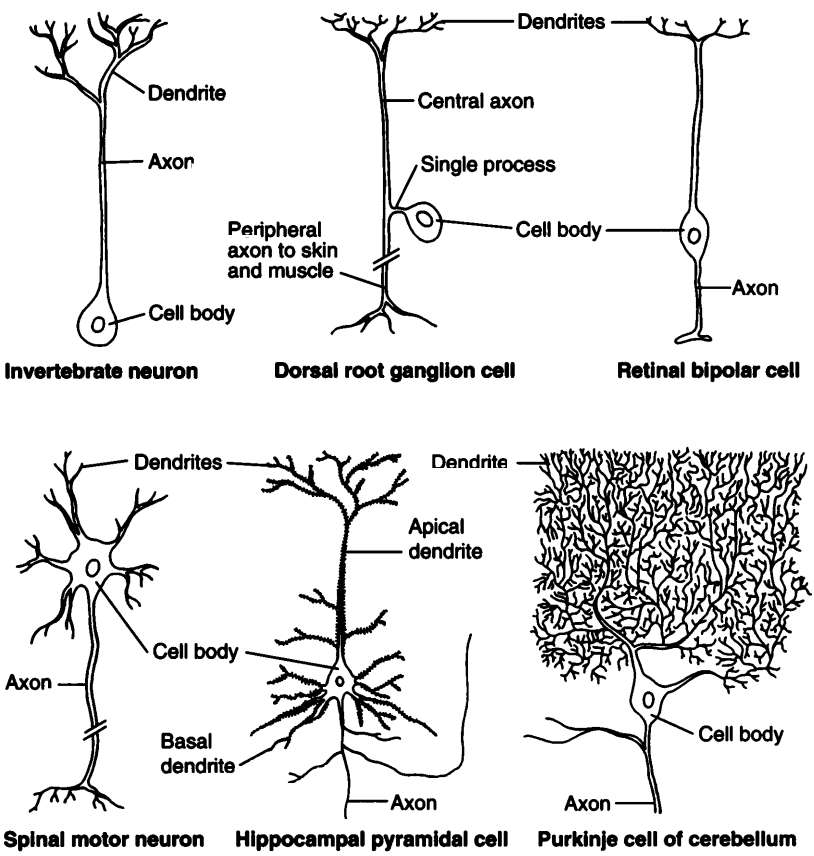

Figure 7.4
Examples of different types of neurons. (From Kandel 1991b.)

shape, whereas the pyramidal cells in the cortex have dendrites that branch less frequently and extend farther (figure 7.4). The dendritic tree of a single neuron receives input from as many as several thousand other neurons through synaptic contacts (which are described below). Input is also received on the soma. In some types of neurons the contact points along the dendrites bulge up, producing little knobs called spines, which change in shape and size with experience, serving to expand the surface area of contact.

Each neuron has only one axon, which has been estimated to make up to 4,000 contacts with the dendrites and somas of as many as 1,000 other neurons (Cherniak 1990; Shepherd 1990b). The axon is connected to the cell body or a large dendrite at a junction called the *axon hillock*. Axons are constructed like fine cylindrical tubes that taper as the tube extends from the soma. At the end of the axon, many small collateral branches extend to make contact with other neurons. As these branches end, the tissue expands into a small bulb known as the *axon terminal*. Some axons are quite short and extend less than a millimeter, whereas others axons, such as those that carry information to or from the spinal cord, can be several feet long.

Neurons are electrically charged, somewhat like a battery, carrying a negative charge inside relative to the outside of their cell membrane. In a resting state the voltage difference or potential between the inside and outside of a neuron is about −70 millivolts. This charge is due to an unequal distribution of charged particles, called ions, between the inside and outside of the cell. Positively charged ions, such as sodium (Na^+), are found in higher concentrations outside the cell, whereas negatively charged ions, such as cell proteins (P^{2-}), are more concentrated inside the cell. Other ions, such as potassium (K^+), calcium (Ca^{2+}), and chloride (Cl^-), are also unequally distributed. This distribution is maintained by the cell membrane, which contains tiny channels, specialized for transporting each type of ion. In its resting state the membrane is semipermeable to K^+, allowing these ions to move relatively easily between the inside and outside of the cell through the potassium channel. However, at rest the membrane is impermeable to Na^+. Cell proteins are too large to pass through the ion channels and remain in the soma. The combined distribution of these ions produces a negatively charged cell.

The cell's charge would eventually run down, were not it for a mechanism in the membrane, called the *sodium-potassium pump*, which pumps potassium back into the cell and sodium out. The cell's negative potential is maintained primarily by keeping Na^+ ions outside of the cell, as potassium channels allow K^+ to leak back into the cell. Ion channels also are activated by changes in the potential of the cell.

Local and Action Potentials Within neurons information is conducted in two different ways. Long-distance transmission occurs via *action potentials*, which travel along the surface of the axon from the soma to the axon terminals. To assimilate incoming information in the dendrites and soma, neurons use *local potentials*. Sometimes local potentials also function to transmit information over very short distances (less than a micrometer) between two adjacent neurons. Action potentials produce a large, all-or-none signal that is brief (1–10 milliseconds) and travels unattenuated and at a rapid rate (as high as 100 meters per second). Local potentials are small and graded, and propagate passively, rapidly degrading as electrical resistance is encountered.

Changes in local potentials of dendrites and the soma are produced by synaptic contact with other neurons (a process described below). These synapses either

electrically excite or inhibit the receiving neuron. The charge of the soma can be depolarized (say, from −70 millivolts to −60 millivolts) if the sum of the local potentials synaptically generated by other neurons is large enough. The more excitatory synapses that occur at the same time, the larger the local potential.

Action potentials begin at the axon hillock. The hillock is very sensitive to changes in the membrane potential of the cell. When the potential in the hillock is depolarized to a critical threshold, the ion channels, which are sensitive to voltage changes, open briefly (about half a millisecond) to allow Na^+ to rush into the cell. Because the concentration of Na^+ outside the cell is much higher than inside, Na^+ is driven inside the cell by processes that tend to equalize the concentrations (diffusion, due to the concentration gradient, and electrostatic pressure, due to the charge gradient). This influx of positive ions rapidly depolarizes the cell further, achieving a positive 50-millivolt charge, at which point the sodium channels close (see figure 7.5).

Meanwhile the potassium channels are also opened to allow K^+ to flow out of the cell through the diffusion process. Because potassium channels open more slowly and the concentration gradient for sodium is more out of balance, the Na^+ influx dominates

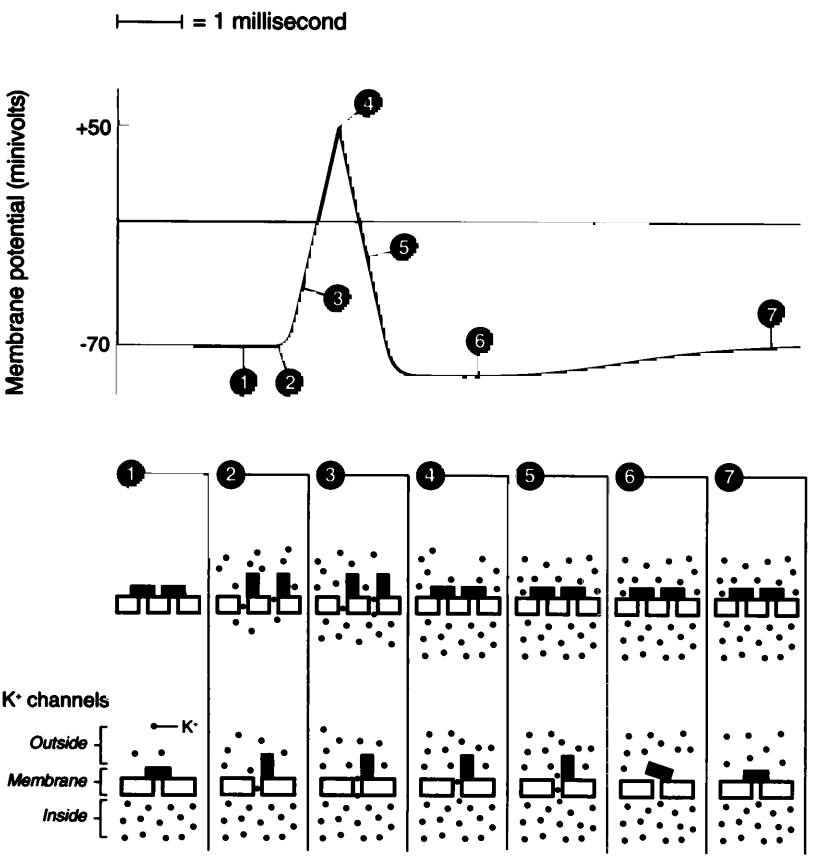

Figure 7.5
Schematic diagram of changes that occur in a neuron's membrane potential during an action potential. Refer to the text for an explanation. (From Thompson 1994.)

the initial changes in the cell's membrane potential. However, after the sodium channels close, the potassium channels remain open. The cell is then negatively repolarized as K^+ continues to leave the cell and as Na^+ is actively transported out of the cell by the sodium-potassium pump. When the cell begins to reach its resting potential, the potassium channels close. However, this process takes a few milliseconds, during which the membrane is actually hyperpolarized (at, say, −75 millivolts) until the sodium-potassium pump and the leaking of Na^+ to the outside of the cell can return the membrane potential to its resting level.

When the hillock begins to depolarize, the adjacent region of axon membrane becomes depolarized as well by the influx of Na^+. These positive ions make the inside of the cell less negatively charged. This reduction in membrane potential is sufficient to trigger adjacent sodium channels, causing them to open and thus moving the action potential farther down the axon membrane. When the sodium channels are open, this region of the membrane cannot be electrically stimulated to produce another action potential until the channels close and the membrane begins to repolarize. This portion of the action potential is called the *refractory period*. Because of the refractory period, the action potential is propagated in only one direction, away from the soma, advancing along that part of the axon membrane which is electrically at rest. Each step in this process is illustrated in figure 7.5.

To speed up the rate at which action potentials are conducted, some axons are sheathed with *Schwann cells*, a type of glial cell. Schwann cells produce a fatty substance called *myelin* that wraps around the axon in multiple layers. Myelin is an electrical insulator that isolates the axon membrane from the extracellular fluid and thus from changes in electrical potential caused by the movement of ions across the membrane. If the myelin sheath were continuous along the axon, the action potential could not propagate. However, every millimeter or so there are small gaps in the myelin, called *nodes of Ranvier*. The electrical current associated with the influx of Na^+ at one node is directly and very rapidly conducted to the next node, causing a depolarization at that node. In this way the action potential jumps from one node to the next. In unmyelinated axons, where the action potential must be propagated continuously along the membrane, conduction speed is rarely more than a few meters per second, but myelination can increase the speed to up to 100 meters per second.

Human infants are born with minimal myelination, mainly in primary sensory and motor areas, but myelination continues throughout childhood, past the age of fifteen. The myelination process has a profound effect on the development of cognitive and motor functions. For example, infants begin to walk at about the time that the peripheral motor neurons that control leg muscles complete myelination. Some researchers have speculated that myelination in associative and frontal cortex may play an important functional role in the development of language and other higher cognitive functions (Huttenlocher 1979). Multiple sclerosis (MS) is a degenerative disease in which patches of myelin and sometimes axons in the motor and sensory tracts are destroyed.

In addition to conducting action potentials, axons provide a conduit for the transportation of chemical substances between the soma and the axon terminal. Certain chemicals, called *neurotransmitters*, used in the synaptic transmission at the axon terminal (described below) are manufactured in the soma and then go through anterograde (forward) transport to the axon terminal, where they are stored in synaptic vesicles. These chemicals also undergo retrograde (backward) transport from the terminal to the soma so they can be reused. A slow (one millimeter per day) anterograde transport

system carries material important for cell growth and regeneration. A faster (10–20 millimeters per day) system carries the transmitter material. Some neurotransmitters (such as acetylcholine, which is involved in synaptic transmission in the peripheral nervous system to active muscle tissue) are manufactured at the axon terminal.

The firing rate of a neuron can range from a few action potentials per second to several hundred depending upon the amount of excitatory or inhibitory input. Changes in the firing rate represent a basic computational property of the nervous system. When a neuron changes its firing rate in response to a particular stimulus characteristic, the neuron is probably involved in the computational analysis of that attribute.

Synapses *Synapses* are the basic building blocks for neural computation. A synapse is a site at which electrical or chemical transmission occurs between neurons. The changes that occur in the firing rates of neurons are controlled by the synaptic activity. Most synaptic transmission in the mammalian brain is chemical. The sending (or presynaptic) cell releases a neurotransmitter, which binds to the membrane surface of the receiving (or postsynaptic) cell, causing a change in its local potential. Some transmissions are *excitatory* and depolarize the postsynaptic cell, driving it toward its firing threshold and thus making it more likely that an action potential will be generated. Other transmissions are *inhibitory* and hyperpolarize a cell, driving the membrane potential below its resting level, reducing the likelihood that an action potential will be produced.

The process of chemical transmission between neurons is called *exocytosis* (out of the cell). Exocytosis takes less than a millisecond. A neurotransmitter is contained in tiny vesicles in the presynaptic terminal. Each vesicle contains about 10,000 molecules of the transmitter chemical. When an action potential reaches the axon terminal, the change in voltage triggers the opening of calcium ion (Ca^{2+}) channels. These ions diffuse into the cell, further depolarizing the terminal and causing the vesicles of neurotransmitter to fuse with the terminal membrane and open to the outside of the cell, releasing the transmitter chemical into a very small space roughly 20 nanometers wide (called the *synaptic cleft*) between the pre- and postsynaptic membranes. The vesicles are thought to be reformed out of the terminal membrane.

The Fast Synaptic System Once released into the synaptic cleft, the neurotransmitter diffuses across the space and attaches to molecular receptors on the postsynaptic membrane (see figure 7.6). These receptors are large protein molecules embedded in the cell membrane with one surface area sticking out. This surface has a region with a precise shape that matches the configuration of the transmitter molecule so that the transmitter can attach itself to this region like a key fitting a lock. In some receptors the attachment of the transmitter alters the receptor's molecular structure so that a channel is opened to the outside of the cell, allowing certain types of ions to diffuse into or out of the postsynaptic cell. Whereas the ion channels involved in the transmission of the action potential are electrically gated—that is, they open or close to changes in the voltage of the cell—these postsynaptic receptors are chemically gated, so that ion channels open and close in the presence of certain neurotransmitter chemicals.

Whether the postsynaptic cell is excited or inhibited depends upon the flow of ions. Excitation causes what is called an *excitatory postsynaptic potential* (EPSP). An EPSP is a depolarization of the cell membrane, driving the membrane potential in a positive direction toward its firing threshold (see the area marked 3 in figure 7.5). An inhibitory

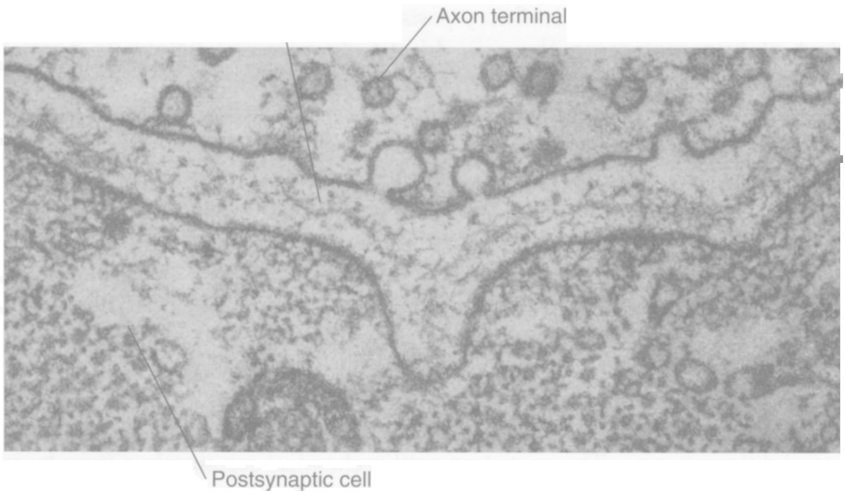

(*a*)

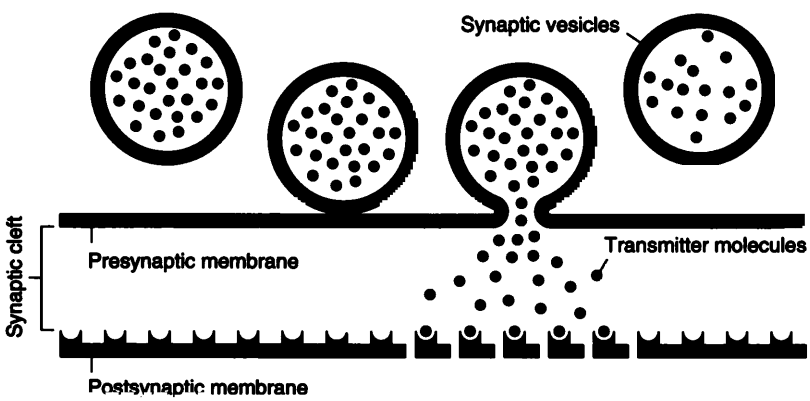

(*b*)

Figure 7.6
(*a*) shows an electron microscope photograph of a synapse. (*b*) is a schematic diagram of this process, showing a synaptic vesicle fusing with the presynaptic membrane to release neurotransmitter into the synaptic cleft. (From Steven 1979.)

synaptic transmission produces an *inhibitory postsynaptic potential* (IPSP) and has the opposite effect of hyperpolarizing the cell, driving its potential below its normal resting level. Acetylcholine, a neurotransmitter that occurs in synapses between motor neurons and muscle cells in the peripheral nervous system, is excitatory because Na^+ channels open at the receptor site and allow positively charged ions to flow into the cell and depolarize it. On the other hand, gamma-aminobutyric acid (GABA) is an inhibitory neurotransmitter because its receptors open Cl^- ion channels, allowing ions to flow into the cell and causing hyperpolarization via their negative charge. In general, the action of the neurotransmitter depends upon the receptor molecules. The same neurotransmitter can be excitatory at some synapses and inhibitory at others; for example, acetylcholine inhibits heart muscle fibers.

Neurotransmitter chemicals are rapidly inactivated so that the timing of the signal is precisely controlled; otherwise, postsynaptic neurons would continue to respond until the brain went out of control (something like an epileptic seizure). Inactivation occurs in several different ways. The neurotransmitter chemical can simply diffuse from the synaptic junction into intercellular space, enzymes in the synaptic cleft can bind to the neurotransmitter and inactivate it, or the transmitter chemicals can be recycled back into the presynaptic terminal and reused, a process called *pinocytosis.*

Second Messenger Systems Other neurotransmitters, such as dopamine and norepinephrine, operate by different and more elaborate mechanisms than chemically gated ion channels. These neurotransmitters affect the concentrations of certain chemical substances in the postsynaptic cells, which in turn cause a chain of chemical events that eventually act on the ion channels directly. This type of postsynaptic effect is termed a second messenger system.

Second messenger systems are extremely complicated and involve a sequence of many chemical events, each of which has multiple effects on other chemical processes in the postsynaptic cell (see figure 7.7). For example, many neurotransmitters use cyclic adenosine monophosphate (cAMP) as a secondary chemical messenger. The precursor to cAMP is adenosine triphosphate (ATP), which is a source of energy that is formed by the metabolism of blood glucose. When the transmitter binds to the postsynaptic receptor, it activates an enzyme, adenylate cyclase, which catalyzes the conversion of ATP into cAMP. In turn, cAMP can act on other biochemical machinery, such as protein kinase (an enzyme that alters the permeability of ion channels), the sodium-potassium pump, the genetic activity of the cell nucleus, and so on. Thus, information from the neurotransmitter, or the first messenger, is multiplied and amplified as the message is passed from one chemical reaction to another. The net result is that a weak signal from the transmitter receptor can produce large and long-lasting effects on the postsynaptic potential. This system is relatively slow. Whereas the chemically gated transmission takes less than a millisecond, second messenger systems require several milliseconds, and some chemical responses can take up to several minutes to be completed.

Some neurotransmitters function in both the fast synaptic and the second messenger systems. For example, in the peripheral nervous system acetylcholine receptors use the fast synaptic system mediated by sodium ion channels; however, in the brain and spinal cord acetylcholine's effects are produced through a second messenger system. Dopamine also uses at least three different types of receptors. The first dopamine receptor (D1) activates a second messenger system, whereas D2 activates chemically gated ion channels.

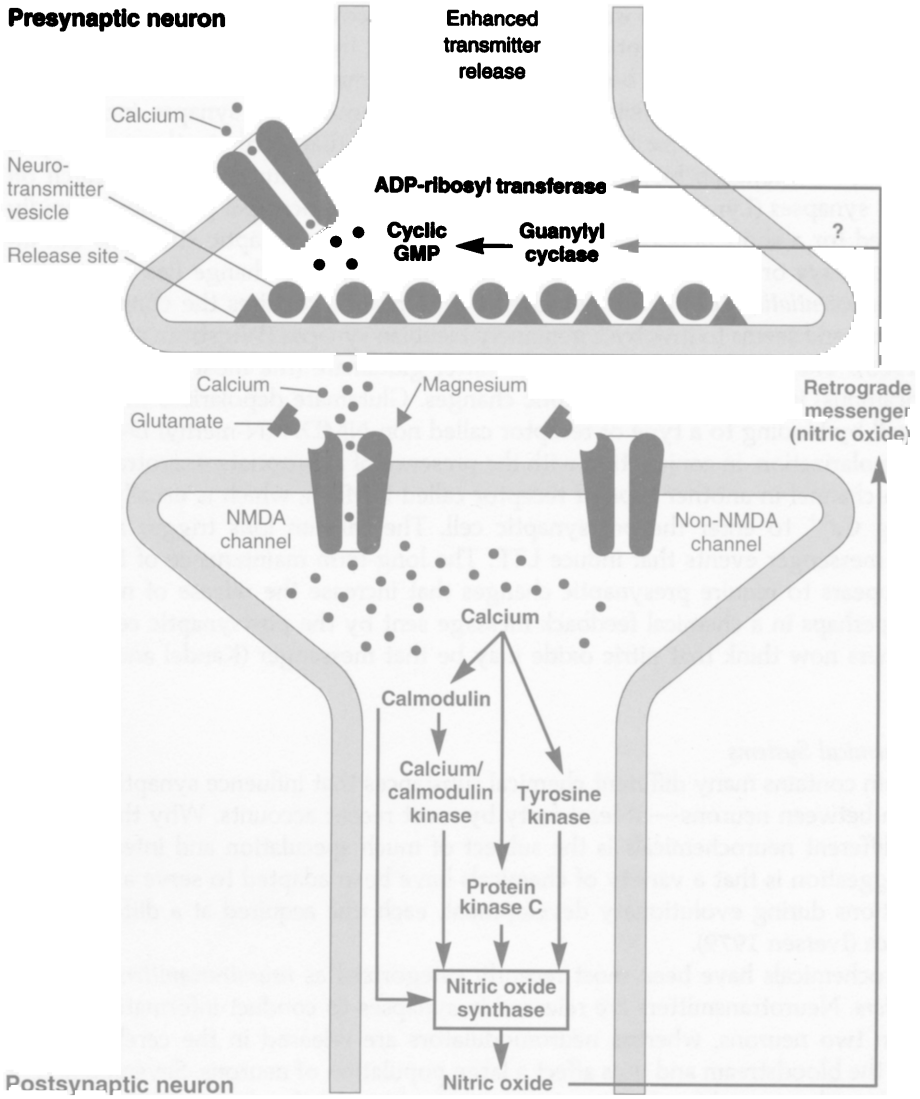

Figure 7.7
Mode of long-term potentiation. Refer to the text for an explanation. (From Kandel and Hawkins 1993.)

Synaptic Modulatory Mechanisms Synaptic effects are modulated by many different mechanisms, which control the size and the duration of the EPSP or IPSP. At least two different mechanisms are known. First, Kandel and his colleagues (Hawkins et al. 1983) reported that the strength of a synaptic connection could be modulated by the activity of a third neuron, which was active at the same time as the presynaptic neuron. This modulatory effect was first found in the study of the classically conditioned gill-withdrawal reflex in a sea snail (aplysia), but it has not yet been seen in the mammalian nervous system. In the aplysia the modulatory neuron stimulates the presynaptic neuron through a serotonin receptor, which leads to activity in the presynaptic neuron that increases the influx of Ca^{2+}. This increased level of calcium ions, in turn, results in more neurotransmitter being released, producing a larger EPSP or IPSP.

Second, Hebb (1949) proposed that coincidental activity in pre- and postsynaptic neurons strengthens the synaptic connection between them. In a classic Hebbian synapse changes are produced in both the pre- and postsynaptic cells, making the connection between them more efficient, whereas in the modulatory synapse increases in synaptic strength are primarily the result of changes that occur in the presynaptic neuron. The mammalian hippocampus has been the site of much recent research on Hebbian synapses (Lynch and Baudry 1984). When a hippocampal pathway is rapidly stimulated for a short period, producing action potentials, synaptic strengths can increase for days or even weeks (Bliss and Lomo 1973). This change has been called *long-term potentiation* (LTP). The associative type of LTP requires the conjunction of two inputs and seems to involve a genuinely Hebbian synapse (Wigstrom and Gustafsson 1985). The release of the neurotransmitter glutamate (the most common in the hippocampus) initiates the postsynaptic changes. Glutamate depolarizes the postsynaptic cell by binding to a type of receptor called non-NMDA (N-methyl D-aspartate). This depolarization, in conjunction with the presence of appropriate neurotransmitters, opens a channel in another type of receptor called NMDA, which is usually blocked, allowing Ca^{2+} to enter the postsynaptic cell. The calcium ions trigger a chain of second messenger events that induce LTP. The long-term maintenance of LTP, however, appears to require presynaptic changes that increase the release of neurotransmitter, perhaps in a chemical feedback message sent by the postsynaptic cell. Several researchers now think that nitric oxide may be that messenger (Kandel and Hawkins 1993).

Neurochemical Systems

The brain contains many different chemical substances that influence synaptic communication between neurons—at least forty by most recent accounts. Why there are so many different neurochemicals is the subject of much speculation and intense study. One suggestion is that a variety of chemicals have been adapted to serve a multitude of functions during evolutionary development, each one acquired at a different time and place (Iversen 1979).

Neurochemicals have been most recently categorized as *neurotransmitters* or *neuromodulators*. Neurotransmitters are released in synapses to conduct information locally between two neurons, whereas neuromodulators are released in the cerebral spinal fluid or the bloodstream and thus affect a large population of neurons. Several different criteria have been used for defining a neurotransmitter: (1) the chemical is synthesized in the presynaptic neuron; (2) it is released from the synaptic terminal; (3) it produces an EPSP or IPSP; and (4) it is removed from the synaptic cleft or deactivated (Feldman and Quenzer 1984). At present only nine neurochemicals have received widespread acceptance as neurotransmitters (Pinel 1990). Neurotransmitters were once thought to have specific anatomical pathways of their own in the brain. Neurons were identified by the neurotransmitter they released, and imaging techniques were employed to produce neuroanatomical maps of the different neurotransmitter systems. However, new evidence demonstrates that some neurons secrete more than one neurotransmitter or have receptors that respond to different types of neuromodulators and neurotransmitters (Snyder 1986). Furthermore, assigning functional specificity to the neurochemical pathways remains elusive. In addition to their overlapping anatomical locations, neurotransmitters differ in their modes of action. Some neurotransmitters modulate the effects of another transmitter on a postsynaptic cell. Some have their

effects through postsynaptic chemical receptors, and others act to modulate voltage-sensitive channels in presynaptic terminals. Still others modify second messenger systems in postsynaptic cells. We are just beginning to understand the full range of neurotransmitter function.

Neuromodulators act more globally, yet appear to have behavioral effects that are highly specific. They can be found in the central nervous system as well as in other areas of the body, such as the gastrointestinal tract (e.g., substance P, enkephalins, somatostatin). Very small amounts of some neuromodulators have been shown to produce profound behavioral effects. For example, an injection of a nanogram (one billionth of a gram) of the peptide angiotensin II can produce intense and prolonged drinking behavior in animals that were not thirsty. Some neuroscientists have interpreted these data to suggest that neuromodulators are specialized for triggering brain activity associated with particular behavioral functions, such as regulating emotional states or the balance of body fluids.

Neurochemicals that exist naturally in the body are called *endogenous* and are distinguished from *exogenous* substances, such as drugs, which can produce profound psychological and neurophysiological effects. Some drugs are so potent that neuroscientists began searching for endogenous substances with the hypothesis that if drugs are so strong, there must be endogenous substances in the brain that have similar neurophysiological properties. Hughes and Kosterlitz (Hughes et al. 1975) discovered two endogenous neurochemicals that acted like morphine, a narcotic drug that is highly addictive and a powerful analgesic for pain. They called these substances enkephalins. Snyder and his colleagues (Snyder 1980) injected radioactively labeled morphine into a laboratory animal and then studied where the isotope would be found in the animal's central nervous system, photographing slices of the brain using film that was sensitive to the presence of the isotope. They discovered that certain regions of the brain and spinal cord contain cells with chemical receptors that have a specific affinity for opiates (so named for the opium poppy from which morphine is derived). Hughes and Kosterlitz showed that enkephalins bound to the opiate receptors just like morphine.

Psychoactive and Neuroleptic Drugs Many other *psychoactive* drugs may produce their effects by mimicking, enhancing, or disrupting the effects of endogenous neurochemicals. For example, several hallucinogens have molecular structures that resemble those of neurotransmitters. LSD is structurally similar to serotonin, and mescaline resembles norepinephrine and dopamine. Caffeine inhibits the enzyme that degrades cAMP in the second messenger system of postsynaptic cells. Cocaine blocks the reuptake of norepinephrine, with the result that more neurotransmitter remains in the synaptic cleft to stimulate receptors. This information does not explain why caffeine acts as a mild psychological stimulant, why LSD distorts sensory perceptions, or why cocaine produces a euphoric effect. However, understanding the behavioral effects of these drugs must certainly include a clear picture of their neuropharmacological properties.

Neuroleptic drugs, so called for their clinical effect on brain function, have received intensive scrutiny in an effort to better understand the biological basis for the mental disorders they treat. Many of these drugs were prescribed over the years because they were clinically effective, but no one knew why they worked. Neuroscientists have been studying how they affect brain function in the hope of finding evidence about the biological bases of mental disorders.

The best example of this approach has been the study of schizophrenia, a mental disorder widely treated with a class of drugs called phenothiazines. Research now indicates that phenothiazines block the receptor sites for the neurotransmitter dopamine, specifically, the third type of dopamine receptor (D3), thus forming the basis for the dopaminergic hypothesis for schizophrenia: if phenothiazines improve schizophrenic symptoms and also block dopamine receptors, perhaps the symptoms are a product of too much dopaminergic activity in the central nervous system (Snyder 1976). Of course, the story is not necessarily simple. One or more of a number of presynaptic and postsynaptic processes could be involved in schizophrenia. Because of the complicated nature of the synaptic process, it is difficult to pin down exactly what the problem is. Additionally, psychological factors, such as stress, might alter neurochemical functions, further complicating the picture. Still, results from this line of research are promising. Other disorders, including depression, mania, obsessive-compulsive disorders, Alzheimer's disease, and Parkinson's disease, have also been the subject of extensive neurochemical study. The potential to alleviate many of the debilitating effects of brain disease by neurochemical means still holds great promise for the future.

Neural Development

How the brain grows from a single cell to a complex structure of billions of neurons largely remains a mystery. During gestation the brain grows at an average rate of several hundred thousand neurons per minute. Animal research has revealed much about the sequence of developmental events, but most of the finer details about what controls and directs the events remain to be discovered. By roughly three weeks after conception, the dorsal surface of the developing embryo contains a small patch of cells that will eventually form the entire human nervous system. Over the next week this plate of approximately 125,000 cells forms a groove that folds over on itself and then fuses to become a tube. This tube eventually becomes the spinal canal and ventricles. Once the neural tube is formed, cells rapidly proliferate, and after forty days of gestation the tube has developed three bulges that eventually become the forebrain, midbrain, and hindbrain (see figure 7.8).

Cells undergo a fixed sequence of events during development: *proliferation, migration,* and *aggregation.* The timing of these processes is critical and varies for different species and for different parts of the nervous system. First, cells divide along the inner wall of the neural tube in a region called the ventricular zone. Each neuron passes through a number of divisions (mitosis), eventually losing the capacity for further division. Neurons stop dividing when they begin migration. Proliferation continues in some parts of the nervous system until a gestational age of about twenty-eight weeks for humans.

Once mitosis is over, cells begin a migration away from the ventricular zone toward the outside wall of the neural tube, forming progressively thicker layers of cells. Generally, larger neurons with long axonal projections begin migration before smaller neurons. Migrating cells move along glial cells, which set up long fibers that act like scaffolds extending from the ventricular zone to the outside wall of the neural tube. Neurons follow these glial processes at a slow rate, on the average of a tenth of a millimeter per day. Cells that eventually occupy the same cortical location always begin migration together.

When neurons reach their definitive "addresses," they aggregate with other neurons to form cortical layers or nuclei. Cells of the same kind tend to group together and

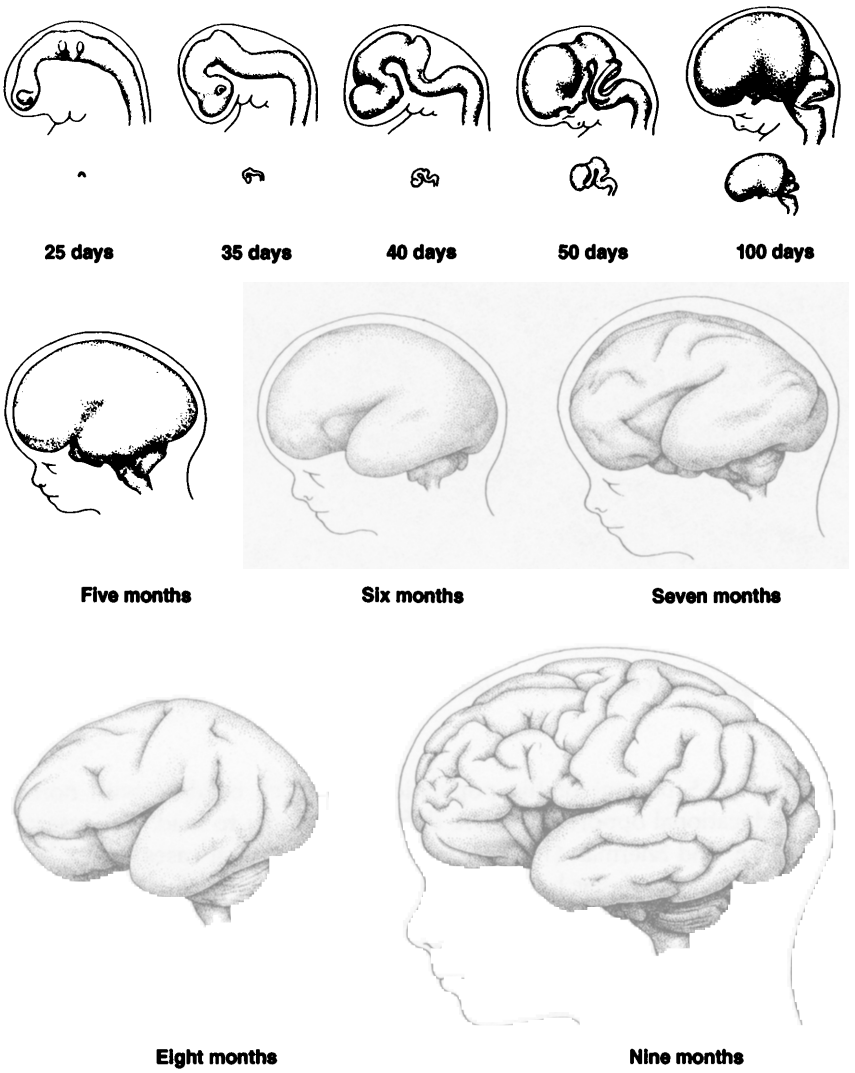

Figure 7.8
Human brain development from 25 days gestation until birth. Figures are enlarged for ages 25 to 100 days (approximate size displayed below). (From Cowan 1979.)

align themselves in a preferred orientation; for example, in the cortex pyramidal neurons are oriented with their dendrites spread out toward the surface of the brain and their axons projecting below. This aggregation process is now thought to be mediated by certain molecules on the surface of the neuron (Rutishauser et al. 1988).

Cell migration errors are not unusual. Through staining and microscopic study, aggregates of neurons can be found in the wrong cortical layer, an abnormality called an *ectopia*. Another migration error called *dysplasia* occurs when a pocket of neural tissue is missing cells (see figure 7.9). In one animal study about 3 percent of the neurons out of the population examined migrated to the wrong location (Cowan 1979).

Several human developmental disorders are associated with cell migration errors. Probably the most carefully studied example is developmental dyslexia, a fairly

Figure 7.9
An example of an ectopia (arrow) and a dysplasia (arrowhead). Bar in lower right corner equals 500 μ. (From Galaburda et al. 1985.)

frequent condition (2–5 percent of the population) in which children with normal intelligence and educational opportunity have difficulty learning to read. The research of Galaburda, Rosen, and Sherman (1989) has documented several cases where large numbers of ectopias and dysplasias have been found in the brains of dyslexic patients, which were studied after the patients' deaths. It has not yet been determined whether cell migration errors are present in all cases of dyslexia or whether such errors actually cause reading problems.

This example suggests the possibility that some developmental disorders may be the product of disruptive events that occurred during critical phases of neural development. Because different parts of the brain are at different stages of development during gestation, the timing of a disruption is critical and controls where the abnormality will appear in the brain and the severity of the abnormality. What could cause such disruptions is the subject of considerable speculation; environmental factors, such as viruses, and genetic defects have both been suggested.

Once migration and aggregation are completed, cells undergo a process of differentiation in which axons grow, synaptic connections are formed, and circuits are pruned down to functional pathways. The early stages of the differentiation process are under genetic control. Studies of neural development in simpler invertebrate nervous systems have shown that for each member of a species, growing neurons follow specific paths and form precise connections with other neurons (Bastiani et al. 1985). Three different hypotheses have been proposed to explain how this takes place: (1) that chemical signals are transmitted between the pre- and postsynaptic cells to guide the growing axon (Sperry 1963); (2) that the chemical or physical signal is laid down in a pathway for the axon to follow (Singer, Nordlander, and Egar 1979); (3) that the relative

position of the cell bodies in an aggregate of growing neurons influences the direction of axon growth so that the synaptic terminals of the axons end up in the same relative position as their cell bodies (Easter et al. 1985).

Following the growth of axons, the number of synapses and the density of synaptic contacts are significantly reduced. Synaptic reduction is commonly assumed to be caused by competition between neurons for functional connections, although the mechanisms underlying these changes are unknown. One hypothesis suggests that some cells die because they fail to receive a necessary nutrient provided by a successful synaptic contact (Cowan 1979). Studies from different species have suggested that from 15 to 85 percent of the initial neuronal population may die during the period of synaptic consolidation. However, humans may not lose neurons to the same degree as other species. Recent studies of human brains have shown a decrease in synaptic contacts without the accompanying cell death (Huttenlocher 1990).

Cell differentiation continues for a long time in humans. Studies of the primary visual cortex have suggested that the density of synaptic connections rapidly increases from about the second month of postnatal life, reaching adult values around the age of five months, and then overproducing for the next three to five months. Connections are then gradually eliminated over the next nine years, reaching adult values by the age of ten (Leuba and Garey 1987). Other parts of the cortex may show a different time course; for example, neuronal density in the frontal lobe does not reach adult values until around seven years of age, and synaptic consolidation may continue through late adolescence (Huttenlocher 1990).

Developmental Plasticity Development in a vertebrate nervous system is marked by progressive phases of overgrowth and subsequent regressive phases in which excess neurons and neural connections are eliminated. Prenatal progressive events include cell proliferation, migration, and aggregation among similar cell types, as described above. Postnatally, cell processes continue to proliferate, producing an abundance of synaptic contacts. Children's synaptic concentrations remain nearly double adult values in some cortical regions until puberty (Huttenlocher 1990). Regressive events begin prenatally with the death of excess neurons, which continues until about the age of two. The pruning of excess synapses and axon collaterals begins postnatally and continues throughout childhood and adolescence (Changeux and Konishi 1987; Cowan et al. 1984; Purves 1988; Rakic 1979).

Developmental changes in cerebral functioning can be studied more directly with PET imaging techniques that measure the rate of cerebral glucose metabolism. The developmental course of the PET data is consistent with neurobiological and microscopic evidence. Cortical metabolic values of newborn babies are 30 percent lower than those of adults, then increase to exceed adult values by the age of two to three years, presumably reflecting the increased metabolic activity necessary to support the many additional synaptic contacts. These PET values remain high until the age of about ten, then gradually decline until they reach adult levels around the age of sixteen to eighteen (Chugani and Phelps 1990).

Regressive events are thought to be the result of competitive interactions between neurons or their processes for some required resource (e.g., a nutritional factor that is necessary for the health of the cell, or electrical stimulation) that is in limited supply (Purves 1988). When a neuron or its processes are deprived of this resource, the size of the axonal and dendritic arbors is reduced, processes are withdrawn from the region

of competition, or the cell degenerates (Changeux and Danchin 1976; Rakic 1986). Through a process of competitive elimination, selective rather than random connections are lost, producing functionally segregated cortical areas. For example, Rakic (1986) has shown that ocular dominance columns in the LGN, which are groups of cells that selectively respond to one eye or the other, are primarily formed by the pruning of axon collaterals from the inappropriate eye.

Early in development there are critical time periods for the establishment of certain behavioral functions. The neural substrates of these behavioral functions develop only if the nervous system receives certain environmental inputs during the critical period. Environmental deprivation can produce profound and permanent impairments, and environmental variation can produce variation in neural development. Critical periods have been demonstrated for a wide variety of behaviors, and in some cases the associated developments in the nervous system have been identified. In the visual system closing one eye during the first few months of life results in more neurons in layer IV of the primary visual cortex being devoted to sensory input from the open eye. The number of layer IV neurons that still can respond to input from the closed eye is correspondingly reduced by 80 percent, thus producing permanent visual damage in that eye (Hubel and Wiesel 1977).

Experiments with white-crowned sparrows and other songbirds have demonstrated that these birds have an innate capacity to develop a basic song pattern specific to their species, but only during the first few months of life (Konishi 1985; Nottebohm 1970). Although the neural correlates have not been identified, an analogous development occurs in human speech. At the age of eight months infants are able to hear the difference between two speech sounds from a foreign language that they have never heard before and that their caretakers cannot distinguish; however, by twelve months of age they lose this ability and can discriminate only speech contrasts from their native language (Werker et al. 1981).

Critical periods for behavioral development correspond to periods during which neural differentiation is occurring in those parts of cortex involved in the behavior. The increase in synaptic density appears to provide a degree of plasticity or flexibility for the formation of behavioral functions. So, for example, because synaptic density continues to remain high in the visual system up to the age of four years, the visual condition known as strabismus (in which the visual axis of one eye tends to deviate from the correct line of sight) can be treated during this period by occluding the sight of the good eye and forcing the use of the squinting eye (Assaf 1982). Presumably a functional plasticity associated with the use of the frontal cortex would persist longer, since synaptic reduction does not begin in this region until about the age of seven.

The plasticity gained from synaptic overproduction also appears to provide the brain some reserve capacity to recover from damage. Many studies have shown age to be a very important factor in successful recovery. In young children, for example, language impairments produced by brain injury are brief, and normal functioning is usually restored when the injury occurs before the age of five, although the severity and location of the injury can impede recovery (Aram, Ekelman, and Gillespie 1990; Kolb and Whishaw 1990).

Since the damaged tissue cannot regenerate itself, linguistic function must have been assimilated into other parts of cortex. The most striking evidence for this conclusion comes from case studies of hemidecorticate patients, who had one cortical hemisphere removed a few weeks after birth because of a rare disease that left only one hemisphere

normal and healthy. Since language functions are localized to the left hemisphere in most people, follow-up studies were remarkable in that many linguistic functions were normal in the patients who had their left hemisphere removed (Dennis and Whitaker 1976). A similar recovery would not occur from a hemispherectomy in adulthood. Although the right hemisphere was not able to process some linguistic tasks well (for example, distinguishing the difference between "the man's lost wallet" and "the lost man's wallet"), patients performed normally on most of the tasks given, including all of the academic tests.

7.3 Neural Representation

The study of the biological properties of the synapse and neuron is an important enterprise, but it cannot by itself tell us how the nervous system represents and transforms meaningful information. To understand how representations and computations are encoded neurally, we must study the properties of interconnected networks of neurons as well. For the cognitive scientist, the study of neural representation and computation is the most exciting, and challenging, frontier of neuroscience.

Structural Principles

Higher cognitive abilities, such as memory consolidation or visual perception, are often associated with large anatomical regions such as the hippocampus or occipital lobe. These regions can be broken down into smaller and smaller structural components, ranging from somewhat smaller anatomically identifiable cortical regions to neural circuits to particular types of neurons or synapses. A full understanding of how the brain achieves some cognitive ability requires understanding the computational functions contributed by each level of structure. For example, the hippocampus can be divided structurally into the dentate gyrus, CA3, and CA1 regions, and different aspects of memory function can be ascribed to these substructures. Rolls (1990) has hypothesized that CA3 serves as an autoassociation system to retrieve the memory of specific autobiographical episodes. Enough is known about the neuroanatomy of CA3 and the response properties of the different types of neurons contained in it that Rolls was able to develop a network model that exhibits some of the required computational properties and is faithful to certain aspects of the neuronal data. Our understanding of the neural implementation of various functions is typically incomplete and still unfolding. Nevertheless, several organizational principles in the nervous system have been well demonstrated. We will touch on a few here; for a more complete discussion, consult some of the readings listed at the end of the chapter.

Experimental Techniques Charting the structural and functional components of the brain has not been easy. Research has relied heavily on three procedures. The first technique involves destroying selected neurons and then studying what other parts of the brain show degeneration as a result of the damage. By tracing the path of degeneration, neuroanatomists hope to discover functional pathways between brain regions. The problem with lesion procedures has been the inability to localize damage to specific cells. Fibers from other cells that may just be passing through the target region also may be destroyed, resulting in their degeneration as well. The recent discovery of acids that selectively destroy neurons with somas in a specific region and not fibers of passage has revived interest in this procedure.

The second technique, which was mentioned earlier in this chapter, involves the introduction of dyes that are taken up by the neural tissue to highlight particular cells or selected parts of them. In the last twenty years new staining procedures have advanced rapidly. One procedure, called *autoradiographic tracing*, uses the anterograde transport system in neurons to determine where a cell's axon terminates. Amino acids are labeled with tritium and injected into the targeted brain site. These acids are taken up by neurons, converted to proteins, and transported down their axons to the synaptic terminal. Sometimes the labeled chemicals are transported to the postsynaptic cells. This transportation takes a few days, after which the animal is killed, and the brain tissue is sliced and prepared in a way that highlights the tritium-stained cells in a photograph. Another procedure uses retrograde tracing to map from synaptic connections back up to the cells from which the axons originate. The most common technique uses horseradish peroxidase (HRP), an enzyme found in horseradish and other plants, which is taken up by the synaptic terminals and transported back up the axon to the cell body. Double labeling, using both techniques, provides a much more precise description of neural pathways than has been possible with older staining or lesion procedures. New techniques for labeling neurotransmitters also have been invented to identify precisely the neurochemical pathways of the brain. Most recently, optical dyes, sensitive to electrical changes or fluctuations in ion concentrations, have been developed that are taken up by synaptically active cells, providing the potential to identify functional neural pathways operating under relatively normal conditions (Grinvald et al. 1986).

A third technique involves recording the electrical activity of single cells. Extracellular recordings are made with microelectrodes, very sharp needles that can be inserted into the brain to monitor the changes in electrical potential of a single neuron. Intracellular recordings are made by micropipettes, very tiny glass tubes that can be inserted into the cell body of a neuron. Although intracellular recordings can only be made for a few minutes, microelectrodes can record for long periods. Using these techniques, neuroscientists have been able to measure the responses of individual neurons to external sensory stimuli. Individual neurons have been shown to increase (or decrease) their firing rates in response to specific features of sensory input, such as the angle of orientation of a visual line display, its direction of movement, or its color (Hubel 1988). However, many neurons that are several synapses beyond primary sensory input do not respond to any simple stimulus characteristic but may be part of an assembly of cells that performs a particular computational function necessary for perception, movement, or thought. Many such neurons receive inputs from multiple sources, suggesting a coordination or integration role. Functions studied at these higher levels include the recognition of complex stimuli, such as faces (Perrett, Mistlin, and Chitty 1987) or the coordination of motor movements for reaching and locomotion (Georgopoulos and Grillner 1989). For a further discussion of these issues, see the section below on computational maps.

Functional Pathways Many different pathways have been discovered in the CNS that carry sensory and motor information. One of the first discoveries was made in 1822 by François Magendie. He reported that the dorsal root of the spinal cord carried sensory information from the peripheral to the central nervous system, but that the ventral root transported motor output to the muscles. This discovery established the principle that different parts of the CNS could be specialized for different functions and that functional pathways could exist to carry specific information.

We know most about the sensory input and motor output systems. Our bodies have five exteroceptive sensory systems, which perceive stimuli from outside the body: vision, touch, hearing, olfaction (smell), and taste. All five sensory systems appear to be organized in a similar fashion. In the classical model, sensory systems are hypothesized to be organized hierarchically, with each major receptor organ passing sensory information along a pathway to the thalamus, which then relays the input to the cortex. Each sensory system has a special part of the thalamus devoted to processing its input. For example, visual information from the retina of the eye is carried to the LGN of the thalamus. The cortex also has primary sensory regions, devoted to sensory input (see figure 7.3). Information is passed from the primary cortex to secondary sensory regions and then to the association cortex, where information from different sensory modalities is integrated. In the simplest model a single stream of information follows a sequential progression, moving from receptors to higher brain regions, gradually being transformed from raw sensory data into perceptual representations. Each level of the system analyzes the information available at that stage and then passes its analysis on to the next level (see figure 7.10).

A hierarchical model of this kind was once proposed for the visual system (Hubel and Wiesel 1977). The model accounts for considerable neurophysiological data. First, there are light-sensitive receptor cells in the retina that signal the intensity of the light at a particular point in the incoming image. The retina's output neurons (called ganglion cells), however, respond to intensity *changes* in small regions of the image. This shift from signaling intensity at a point to signaling local contrast can be seen as the

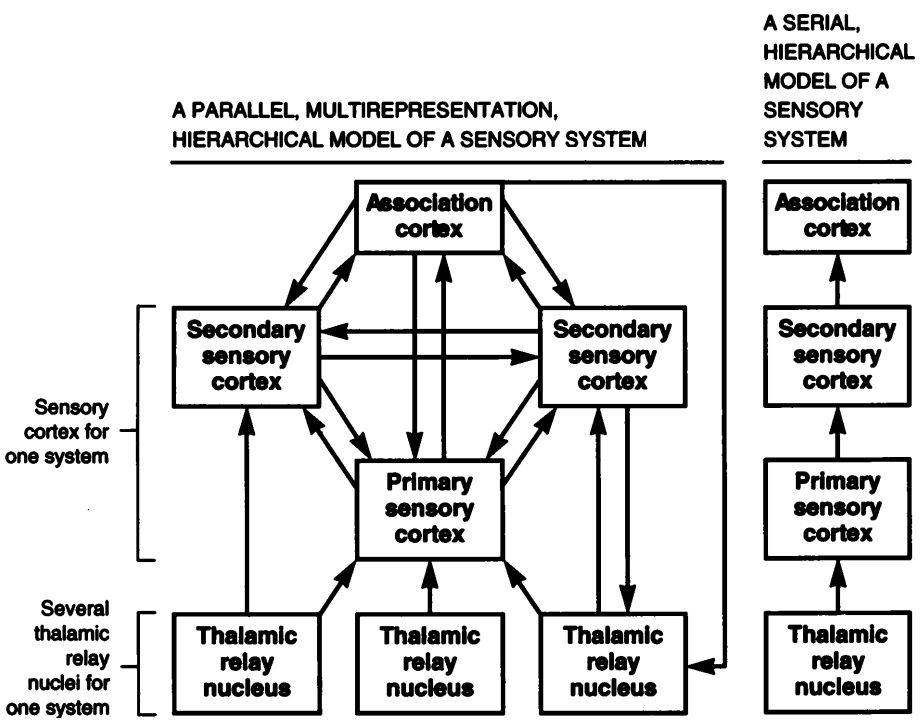

Figure 7.10
Two different models of sensory processing. (Adapted from Pinel 1990.)

first step in a process of object recognition. The regions of local contrast in the image can be seen as early information about the edges or contours of objects. The contrast signals are relayed through the LGN, and the next step in the hierarchy is accomplished by cells in the primary visual cortex (called *simple cells*) that detect contrast boundaries at specific angular orientations in small regions of the image. These *edge* and *line detectors* can be seen as organizing the local, unoriented contrast information coming from the retina into small, oriented segments of the boundaries of objects in the image. Further steps in the process would be associated with other cells in the visual cortex that respond to the presence of an oriented boundary over a larger region of the image (*complex cells*) or that respond to corners or terminated lines (*hypercomplex cells*). Eventually the construction process was hypothesized to reach cells in the medial-temporal region that are specialized to respond to complex shapes, such as a hand or face (Desimone and Gross 1979) or perhaps even specifically to the face of one's grandmother (Barlow 1972).

Although the early evidence on the response properties of neurons in the visual system was consistent with this classical model of a single hierarchical visual pathway, recent evidence has suggested that the system is organized into a number of parallel pathways (see figure 7.10). Information about shape, color, motion, and binocular disparity appears to be processed in distinct parallel channels (Hubel 1988). The parallelism is already present in the ganglion cells. In the primate retina there are two types of ganglion cells, which can be distinguished by the sizes of their cell bodies, the shapes of their dendritic trees, and their response properties. Although the small cells (P-cells) and large cells (M-cells) are intermixed in the retina, their axons project to distinct layers of the LGN. The P-cells project to the *parvocellular* layers and the M-cells to the *magnocellular* layers of the LGN. The separation between the two pathways is maintained, and they are further subdivided, as visual analysis develops in the visual areas of the cortex (Livingstone and Hubel 1987). Computation in the visual system is still thought to be hierarchical in the sense that later stages of the system deliver more reliable and meaningful information about the visual world in comparison with earlier stages, whose outputs code local properties of the visual image that are often not reliably correlated with properties of the visual world (see chapter 12). For example, early stages of the color pathway cannot represent the surface color of an object independently of the color of illuminating light, whereas later stages of the pathway can.

Not much is known about how increasingly meaningful information is developed in the parallel channels of the primate visual system or about how the information is integrated to produce visually guided behavior and object recognition. However, detailed studies of the neural substrates of two perceptually driven behaviors, the sound localization response of the barn owl and the jamming avoidance response of the electric fish (Konishi 1991), have shown that it is possible to analyze a sensory pathway in detail. In both cases the sensory analysis is carried out in two separate parallel channels, which eventually converge on the neurons that code for the critical property.

The analysis of the barn owl begins with the observation that when the owl hears a sound, it turns its head in the direction of the sound source. This head orientation is accurate in both the horizontal (azimuth) and vertical (elevation) directions. Thus, if the source of a sound is 30 degrees to the left and 15 degrees above the axis of the current head orientation, the owl will turn its head upward and to the left the appropriate

amounts. By playing experimentally controlled sounds through earphones, researchers were able to show that the horizontal and vertical components of the orientation response are controlled by two different aspects of the sound. Horizontal localization is based on interaural time differences. Because the owl's ears are separated horizontally in space, a sound in the environment is delivered to the two ears in slightly different ways. Consider a 100-hertz (cycles per second) sound that originates to the left of the owl. The sound will reach the left ear slightly before it reaches the right ear. Another difference arises from the fact that the sound oscillates sinusoidally with time. The sine wave for a 100-hertz tone reaches its peak energy value 100 times per second. Each energy peak will reach the left ear slightly before it reaches the right ear. The sine waves at the two ears are said to be *out of phase* with each other. Complex sounds are mixtures of many frequency (or pitch) components. The experiments showed that the owl uses the phase differences at all audible frequencies rather than the initial difference in onset time for horizontal localization.

Because the owl's right ear is angled slightly upward and the left ear is angled downward, differences in the properties of a sound as it arrives at the two ears can also be used to determine the sound's elevation. The behavioral experiments showed that intensity (or *amplitude*) differences between the two ears were used in this case. A sound that is above the owl is slightly louder in the right ear than in the left and vice versa. As a result of the behavioral studies researchers knew that the auditory system must extract a combination of interaural time and amplitude differences from the incoming sound. This combination specifies the horizontal and vertical values of the sound and hence its direction in space. A code for the combination must be passed on to the visual and motor systems, since the owl turns its head toward sounds in its environment and tries to visually locate the source of the sound.

Konishi and his colleagues (Konishi 1991) were able to locate *space-specific* neurons in the exterior nucleus of the inferior colliculus, a midbrain structure that is part of the auditory system. Single-cell recordings showed that each space-specific neuron fired when a sound occurred in a particular spatial direction. Using experimentally controlled sounds delivered through earphones, the researchers were able to show that the neurons were in fact selective for combinations of interaural time and amplitude differences. The space-specific neurons appear to be the output neurons for the sound localization computation. They directly encode the information that triggers sound localization behavior. Anatomically, they are at the top of the owl's auditory pathway, and their axons project to the optic tectum, a midbrain structure involved in vision containing neurons that respond to both auditory and visual inputs. Electrical stimulation of these auditory-visual neurons causes quick head movements of the kind that are also induced by sound.

The neural circuitry leading up to the space-specific neurons is sketched in figure 7.11. The initial transduction of sound is accomplished by receptor cells in the *cochlea*, a structure in the inner ear. Receptor cells fire in response to particular frequencies. Their rate of firing is correlated with the intensity of the incoming sound, and they tend to fire in phase with the incoming sound. Thus, frequency, amplitude, and phase information are all available at the receptor level. Further amplitude and phase processing, however, occurs in two separate, parallel pathways. The output of each cochlea projects to both a magnocellular nucleus and an angular nucleus. Magnocellular nucleus neurons tend to fire in phase with a particular frequency component of the sound in the source ear. Their responses are *phase-locked* with the incoming stimulus.

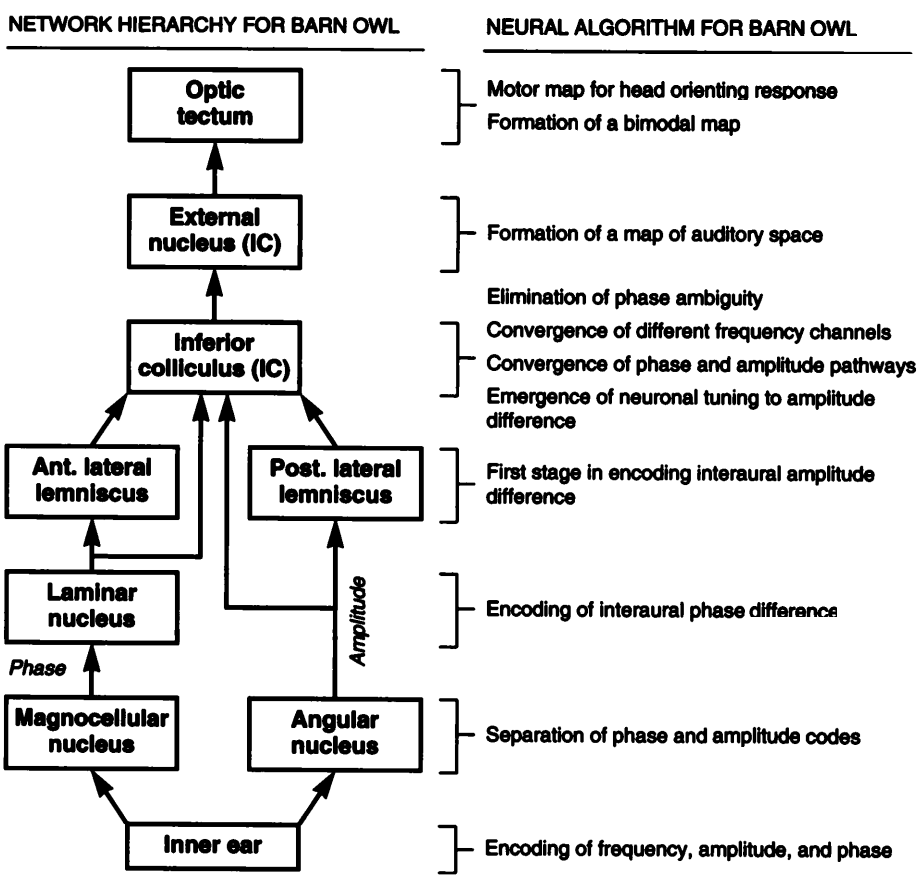

Figure 7.11
An anatomical and functional mapping of the auditory system of the barn owl. (From Konishi 1991.)

These neurons are not sensitive to variations in amplitude, however. In contrast, angular nucleus neurons fire more rapidly

rather than from thalamus to cortex, the architectures depicted in figures 7.10 and 7.11 are similar. The hierarchical nature of the sound localization pathway is very clear. Computationally, the representations at higher levels in the pathway represent behaviorally relevant information more explicitly and unambiguously than the representations at lower levels. Neurons at the top of the hierarchy encode the behaviorally relevant property of spatial direction. Neurons somewhat lower in the hierarchy encode two independent dimensions of spatial direction. Neurons that are still lower in the hierarchy encode stimulus properties that are useful inputs to the computation of the two dimensions, and so on. Anatomically, it is possible to trace the forward-feeding connections among cell populations that are responsible for the levels of representation. In some cases the detailed circuitry of the connections helps explain how a higher-level representation is computed from a lower-level one. For example, the axonal projections from the magnocellular nuclei to the laminar nuclei appear to be organized as delay lines, which allow phase-locked signals that originate at different times at the two ears to arrive simultaneously at laminar neurons.

Sound localization in the barn owl also illustrates other principles of neural coding. One principle is that of *place-coding*. Nervous systems often dedicate specific neurons to the representation of a given type of information. The location of a neuron at a particular place in the neural circuitry and its interconnection with other neurons determine its computational role. A simple example is provided by the fact that cells in the occipital cortex are devoted to vision, whereas cells in the postcentral gyrus are devoted to sensations from the body surface. At a finer grain neurons in a particular region of the visual cortex correspond to a particular region of the retina, and neurons in a particular region of the somatosensory cortex correspond to a particular region of the body surface. In the barn owl we see more subtle examples of this principle. Groups of neurons are dedicated to the representation of more abstract properties, such as monaural phase, monaural amplitude, interaural time differences, interaural amplitude differences, and direction in auditory space. The significance of place-coding is brought out further in the sections below on cortical columns and maps.

In addition to its place in a neural network, a neuron's computational role hinges on details of its moment-to-moment behavior. In the angular nucleus of the barn owl, for example, the amplitude of a sound at a particular frequency is encoded by the rate at which frequency-selective neurons fire. In the magnocellular nucleus, however, phase is encoded by timing action potentials to correspond to the phase of the stimulus. To see that the rate and timing of impulses represent two distinct coding strategies, recall that two neurons could be firing at the same rate but still be out of phase with each other. Rate and phase of firing are well-established bases for neural codes. Other time-varying properties of neuronal activation, such as the probability of firing during a brief interval, may be added to this short list. However, a point made about connectionist models in chapter 2 can be reiterated in a somewhat different form here. It is not thought that the timing of a train of neural impulses can encode arbitrarily complex information in the manner of a text being transmitted over a telegraph line via Morse code.

The study of sound localization in the barn owl is instructive as an example of a research program in computational neuroscience that has advanced to the point where evidence from a cognitive-level analysis of a behavior has converged with evidence from neurophysiology and neuroanatomy to favor a rather detailed model. This is the kind of convergence that we can look forward to in the study of the cognitive

capacities of humans and other primates, but for the most part it is a goal rather than a reality. In judging the current state of cognitive neuroscience, it is a good model to keep in mind. However, the specifics of the example should not be overgeneralized. The single-neuron encoding and the parallel-hierarchical organization of the computation may or may not characterize other capacities in other organisms.

Columnar Organization The primary cortex of all five sensory systems contains vertical columns of cells that tend to respond to the same kinds of sensory input. Each column receives input from the same general area in the peripheral sensory organ, thus encoding the location of sensory input. In the visual system this means that adjacent columns of cells in the occipital lobe correspond to overlapping parts of the visual field. Microcolumns also have been discovered inside the larger place-columns (Hubel 1988). These substructures organize information into left- and right-eye input as well as columns that show a preference for contrast boundaries of particular orientations (see figure 12.10).

Topographic Maps Because spatial relationships between cortical columns correspond to the spatial organization found between adjacent neurons in the sensory organs, primary sensory cortical areas provide detailed representations of sensory input patterns, which are called *topographic maps*. Such maps have been found in the visual, somatosensory, and auditory cortices, but gustatory and olfactory maps have not yet been discovered (however, see Skarda and Freeman 1987 for a discussion of olfactory representation). Each map displays systematic variation for a particular stimulus attribute. In the visual and somatosensory systems, the mappings correspond to the two-dimensional topography of the retina (retinotopic) and body surface (somatotopic), respectively. The primary auditory cortex is mapped tonotopically, with the anterior portion responding to high-frequency tones and posterior regions responding to progressively lower frequencies.

These cortical maps maintain a correspondence with peripheral sensory neurons, so that higher concentrations of sensory cells in one region will produce a central map with more space devoted to that region. Consequently, cortical maps of the hands and face, which have many more sensory neurons than the feet, can appear quite distorted in relation to the size of the body part (see figure 7.12).

Multiple cortical maps appear to exist for each sensory modality in their respective cortical regions. Many retinotopic maps, for example, have been discovered on the cortex of nonhuman primates; the macaque monkey has roughly twenty different visual maps, which are interconnected by over eighty pathways. These maps represent different stimulus attributes, such as the orientation of contrast borders or the direction of movement of visual objects (Swindale, Cynader, and Matsubara 1990). Maps also may perform different computations on the same attribute, or facilitate the speed with which certain computations are performed (Dudai 1989). As researchers learn to tune their experiments to the proper attributes, many more maps probably will be discovered.

Substantial experimental evidence suggests that cortical maps can change with experience and even vary between individuals (Merzenich 1987; Pons, Garraghty, and Mishkin 1988). Merzenich and his colleagues have demonstrated these changes in a most striking way. Through the use of multiple microelectrodes they recorded from different regions in the somatosensory cortex of an owl monkey while stimulating the surface of the monkey's hand and digits. In this way they were able to map the

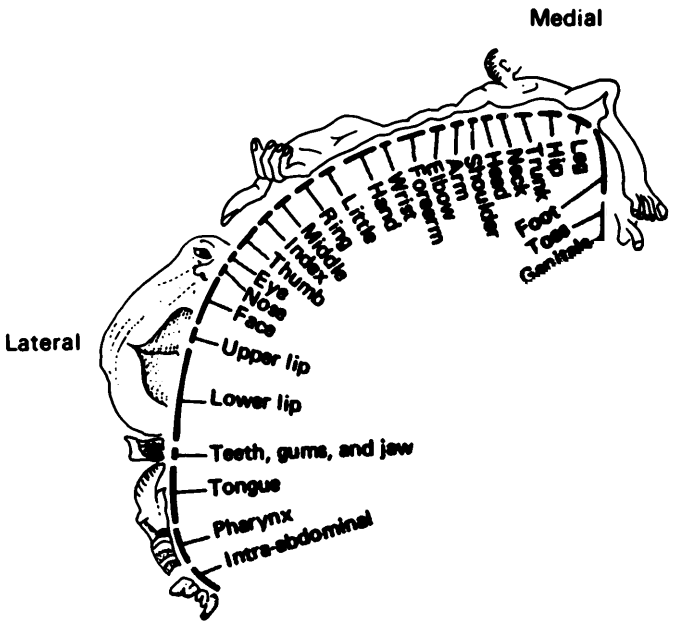

Figure 7.12
Illustration of the primary somatosensory regions corresponding to different parts of the body (From Kandel and Schwartz 1985.)

monkey's hand and digits onto the somatosensory cortex (see figure 7.13). The monkey then received an hour and a half of daily training in which it rotated a disk with its second, third, and occasionally fourth digits in order to receive a food reward. After twenty weeks of training, the researchers remapped the cortex and discovered that the cortical representation of the used digits had markedly expanded. These results suggest a plasticity or ability of the cortex to adapt to experience that previously had not been thought possible.

Cortical maps also vary among individuals. Thirty-five years ago Woolsey (1960) proposed a theory that the auditory cortex was tonotopically organized, with high-frequency tones more anterior than low-frequency tones. Early attempts to test Woolsey's theory suggested he was wrong. When researchers recorded from neurons in a specific region of auditory cortex, they found that cells in different animals responded to a wide range of frequencies rather than to the same frequency, as Woolsey had suggested. When experimental procedures improved, allowing more cells to be recorded from one animal, researchers discovered that different cells from the same region responded to similar frequencies in an individual animal, thus supporting the tonotopic theory (Merzenich, Knight, and Roth 1974). They also found that the part of the auditory cortex that responded to a given frequency varied from one animal to another.

The lesson from this research is clear. Topographical mapping is an important organizing principle that is consistently followed in all sensory modalities studied so far. However, the spatial organization of cortical maps changes with experience, and individual differences are likely to be considerable.

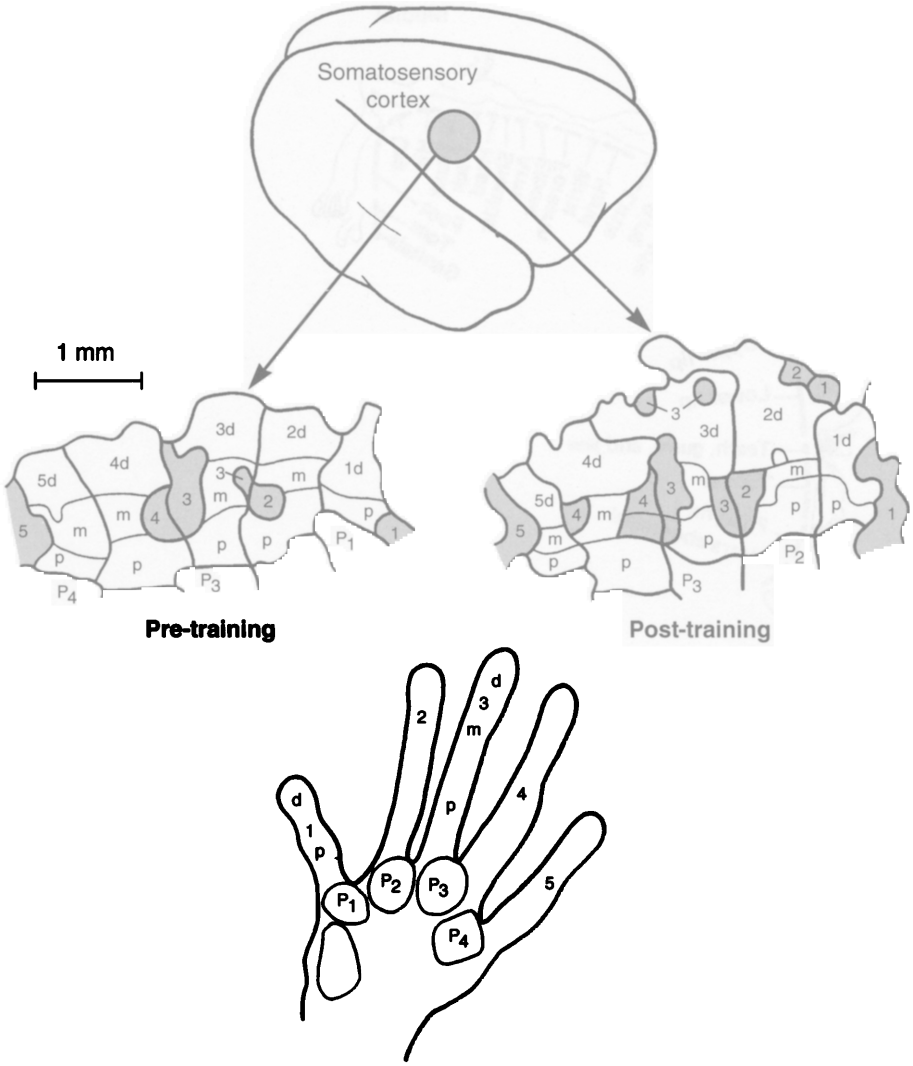

Figure 7.13
A topographical map of the somatosensory cortex of an adult owl monkey's hand. The hand surface of a normal monkey is coded in the bottom figure. Each digit is numbered (thumb = 1); d = distal, m = middle, p = proximal parts of the digits; P1–P4 = the palmar pads at the base of the digits. The surface of the hand is represented in a cortical map in the somatosensory cortex, illustrated in the middle figures. On the left is the cortical map before training. Each cortical area has neurons that respond to touch in the corresponding region of the hand. The gray areas are the dorsal surface of each digit. The right figure shows the same cortical region of the same animal after training. Over twenty weeks, monkeys were trained for one and one-half hours each day to use the distal parts of digits 2 and 3. Notice the increased size of the cortical areas that responded to tactile stimulation in digits 2 and 3. (Adapted from Jenkins and Merzenich 1987.)

Computational Maps The cortex also contains *computational maps* that do not correspond to the topography of sensory neurons. These maps represent more abstract information that figures in cognition or behavior. We saw above, for example, that the barn owl's inferior colliculus (a subcortical structure) contains neurons that code for particular spatial directions. The neurons are arranged in an orderly map, in which neighboring neurons respond to neighboring directions and in which changes in azimuth and elevation are defined along roughly perpendicular axes (Konishi 1986). For echolocation the mustached bat emits a complex spectrum of sounds and then listens to the echoes that return from targets (biosonar). Suga's research (1990) has revealed a central auditory system that not only encodes the frequency, amplitude, and time delay of the echoes but also contains computational maps for target velocity and range.

Maps can also encode combinations of two or more types of information. Topographical information can combine with computed information, within or across sensory modalities (Dudai 1989). In the barn owl we saw that the axons of the space-specific neurons in the inferior colliculus project to the optic tectum, which contains a visual-auditory map of space that responds to both auditory and visual input and that triggers head movements. Swindale and his colleagues (1990) provide an interesting description of how visual orientation and direction maps interact to maintain continuity and completeness in our cortical representations of visual scenes.

Distributed versus Single-cell Representations

We saw in chapter 2 that an issue of single-unit versus distributed coding arises in connectionist modeling. Connectionist models, of course, are typically not specified in any detail at the neural level of analysis. For example, in a network with single-unit coding, each unit might be implemented at the neural level by many neurons in a distributed manner. Nevertheless, we can imagine taking single-unit coding all the way down to the neural level, using single neurons to encode features, concepts, propositions, or schemas. In this case the standard apocryphal example of the grandmother unit would be a single neuron that fired if and only if you were thinking of your grandmother—the so-called grandmother cell. Just as in a connectionist net, such a cell would have to be connected to many other cells, involved in processing visual input, kinship concepts, and so on; nevertheless, activity in this cell would be a necessary and sufficient condition for grandmother cognition.

It is sometimes thought that single-cell coding is unreliable and therefore would be selected against during evolution. Imagine losing your concept of your grandmother when the single cell that encodes her dies randomly as you struggle through the last mile of a marathon. Evolution could guard against this possibility, however, by providing multiple copies of each coding unit. Thus, there might be ten copies of your grandmother cell and a miniscule probability that more than one or two of them would die off in a lifetime, barring a major brain injury or disease. In such a brain, there would be single-unit coding that is resistant to damage.

A theory of single-cell representation is initially attractive conceptually because it is easy to understand how items of information are represented and easy to understand the functional role of the individual neuron in cognition. The theory is also initially attractive methodologically. Since we know how to record the activity of single neurons, we can immediately start looking in experimental animals for neurons whose

activity is uniquely correlated with, say, the recognition of some particular object. This strategy is not so easily implemented, however. For example, to show that a cell is a grandmother cell, we would at least have to show that it fires vigorously to grandmother and not to anything else. But, confronted with a neuron that is fairly well removed from the sensory-motor periphery, we might have to test thousands of objects to find one that it prefers, and if we hit upon a preferred object, we might have to test thousands more to show that the cell does *not* fire to each of those. Further, if we found that the cell responded to some degree to other objects (though not as strongly as to grandmother), we simply would not be able to conclude that the cell does not play a role in the representation of nongrandmothers. So, there is no simple empirical method to test for single-cell representations. We cannot easily confirm or disconfirm the existence of such representations by just wandering through the cortex and testing the response properties of individual neurons. We need a more complete theory of neural representation that makes testable predictions about the anatomy of and the course of processing in entire neural pathways or networks.

Distributed neural encoding is somewhat harder to imagine, but, as we saw in chapter 2, the initial conceptual hurdles are cleared by connectionist models that successfully employ distributed codes in networks of neuronlike units. To repeat, there is no assumption in these models that individual units correspond to single cells at the neural level of analysis, but the models demonstrate that distributed coding is a viable possibility in the brain. Cognition of the proverbial grandmother might now be represented by a distinctive pattern of activity over a population of neurons, all of which are also involved in encoding other concepts. Following the lead of connectionist networks in which knowledge of each piece of information is encoded in the pattern of weights across *all* of the units in the net, we can imagine an extreme case in which knowledge of grandmother (and all other concepts) is encoded in synaptic patterns across all of the trillion or so cells in the cerebral cortex. Currently, this possibility is not being actively pursued by researchers, because, as we have seen, there is abundant evidence that the brain is highly organized structurally. It is unlikely, for example, that the visual recognition of your grandmother involves patterns of activation across the entire primary motor cortex. It is more likely that distributed codes are restricted to functional regions of the brain, such as the cortical maps just discussed.

Like single-cell representations, distributed codes cannot be found through simple fishing expeditions in the brain. Even with the best techniques for simultaneously recording from more than one neuron, for example, it might be difficult to record from enough of the neurons involved in a distributed code to figure out how the code works. If a particular distributed representation depends on only a fraction of the cells in some cortical region, it might also be difficult to find many of those cells. Worse, the fact that a certain profile of activity occurs in a number of cells under some conditions and not in others does not by itself demonstrate that distributed coding is involved. After all, in a network with single-unit coding all of the cells will have some level of activity all of the time, and the profile of activity in some group of cells might well be similar under similar circumstances. For example, activity in the cell that encodes grandmother cognition might well be accompanied by strong activity in the cells that encode gray-hair cognition and twinkling-eye cognition, weak activity in the cell that encodes beard cognition, and middling activity in the cell that encodes a particular nose shape that sort of matches grandmother's. As in the case of single-unit coding,

what is required is a detailed theory and data about anatomical structure and patterns of activity at more than one level in a pathway.

Both local and distributed coding are known to occur in nervous systems. In the barn owl, for example, individual space-specific neurons at the top of the auditory pathway code for particular directions. At lower levels in the hierarchy, however, the coding of spatial direction is distributed between the amplitude and phase pathways. At the inner ear the information about the direction of the auditory stimulus is distributed across virtually all of the receptors in the two ears. Generally, distributed coding is a necessity toward the periphery of sensory systems for at least two reasons. First, cognitively or behaviorally significant information in the stimulus has to be extracted from a multiplicity of often ambiguous cues in the input. For example, single units for sound direction could not exist in or near the inner ear of the barn owl because information from the two ears has to be compared to establish direction. Or, just as obviously, there could not be grandmother detectors on the retina, because when grandmother is in view the information about her is distributed across a fairly large area of the retinal image and therefore across many receptor cells. To recognize grandmother, the outputs of these cells must be combined, and the relevant information coming from them (e.g., hair color) must be separated from the irrelevant (e.g., level of illumination). Second, it is often physically impossible to use single units to encode information that is meaningful and locally available at the periphery. The wavelength of the light at a particular spot on the retinal image is a case in point. Hundreds or thousands of units, each selective for a particular wavelength, could not be packed into each point on the retina. Instead, sensitivity to variations in wavelength is accomplished by just three types of receptors.

The retinal cones are an example of *coarse coding* in the nervous system. Each type of cone responds across a broad region of the visible spectrum. The medium wavelength ("green") cones, for example, have a peak sensitivity around 530 nanometers, but their sensitivity falls off rather gradually around this peak. Because of their broad tuning, the green cones do not represent any particular wavelength. Further, their output is a joint function of their sensitivity and the intensity of the incoming light. Their response to an intense light at, say, 500 nanometers can be stronger than their response to a dim light at 530 nanometers. The result of this trade-off between wavelength and intensity is that a single type of receptor cannot encode wavelength (and therefore color) information at all. However, the combination of the responses of the three broadly tuned cone types (each with a different peak sensitivity) does encode wavelength information. This combination constitutes a distributed code, which can be represented as a vector of three numbers or a point in a three-dimensional space. Even though none of the cone types has the ability to encode wavelength information by itself, the combination gives us the ability to perceive many thousands of colors. Generally, the profile of responses from a population of coarsely tuned units can specify a value with great precision.

The ability of neural systems to develop codes that are more precise than the resolving power of the individual input neurons is known as *hyperacuity*. In addition to color vision, the visual system exhibits a number of other hyperacuities. For example, people can make spatial discriminations involving distances that are smaller than the distance between adjacent receptor cells on the retina. Bats can discriminate distances as small as 1.0 to 1.5 centimeters using the time delays between the return of their

echoes (Suga 1990). The discrimination requires a sensitivity to time differences in the range of 0.06 to 0.09 milliseconds, but the sensitivity of individual auditory neurons is only 0.13 milliseconds. A fertile hypothesis is that distributed coarse codes are a general strategy for achieving hyperacuities (Baldi and Heiligenberg 1988). The research of Georgopoulos and his colleagues (1986) illustrates one approach to pursuing this hypothesis. They recorded the activity of cells in the motor cortex after training monkeys to reach for different spatial locations. Most of the cells showed a response preference for a particular direction of movement, but they were broadly tuned to a range of directions. The response curves of different cells also overlapped considerably. It was impossible to predict the actual direction of movement from the activity of any single neuron. However, the average response of the entire cell population predicted the exact direction of arm movement. Similar results have been obtained for other motor (Schor, Miller, and Tomko 1984) and visual (Steinmetz et al. 1987) behavior.

In the case of arm movement (as well as some other motor behaviors) the response of a cell population represents its overall average response. It is possible that at a higher level of representation the average is represented by finely tuned direction-specific neurons. In other cases, taking the average response over the population would destroy the information it encodes. For example, averaging the responses of the three retinal cone types would produce a single broadly tuned response curve similar to that of each of the individual cones, with no capacity to encode wavelength information. In these cases any higher-order grandmother units would have to capture properties of the profile of responses over the population. Alternatively, vectors of activation can be mapped directly onto other vectors with no intervening local code.

The hypothesis that distributed coding is employed very widely by the nervous system currently guides a great deal of work in computational neuroscience (Churchland and Sejnowski 1992). The hypothesis is supported by the known cases of distributed coding and by the finding that neurons in sensory and motor cortical areas tend to be broadly tuned. It is also supported by the necessity that even where local coding holds at the top of a processing hierarchy, there must be distributed coding at lower levels. Where there is a grandmother cell, there must be distributed information about grandmother below the level of that cell. The hypothesis is also fueled by the theoretical arguments for distributed coding that have emerged from connectionist research (Hinton, McClelland, and Rumelhart 1986, and chapter 2). With distributed coding, similar objects or concepts have similar codes, sharing many elements or having similar values on many features. The system can respond to a novel input vector by exploiting its correlations with familiar vectors. It can map incomplete or noisy inputs to decent outputs. The high-dimensional spaces created by distributed representations contain plenty of room for encoding vast numbers of objects or concepts and their variants, produced by different points of view or different interpretations. Adjustable weights allow distributed systems to learn complex mappings from experience. Such systems continue to respond reasonably even when their connections and weights have sustained significant damage. These properties, which seem to be properties of nervous systems, are a natural consequence of distributed coding, and it is often argued that it is not clear how to achieve them in systems that rely extensively on local coding. The strongest form of the theoretical argument is that distributed codes should be maintained throughout a system, mapping vectors onto vectors, without passing through stages of local coding. Much of the evidence is consistent with this view, but it is

consistent as well with a view that gives a strong role to local coding at higher levels of representation, which also has its theoretical proponents (Trehub 1991). The general debate about strategies of neural coding and computation can only be settled by detailed models and empirical evidence.

Innate or Learned Representations

Philosophers and scientists have proposed conflicting views on exactly how representations are formed in a developing nervous system. In general terms the empiricist position argues that neural representations are a reflection of environmental stimulation. Neural circuits are programmed as if the brain were a blank slate to be written on by experience. The opposing nativist position hypothesizes that representations are innate and biologically "hardwired." Neural circuits are already in place to subserve specialized functions, and the "right" kind of sensory input triggers the activation of such prewired circuits.

Neurodevelopmental and genetic considerations suggest that both positions have merit. The major neural pathways and structures are very similar in all mammals, suggesting a high degree of genetic hardwiring to guide the design. Members of the same species produce nearly identical neural circuits during the early stages of neural development. However, Changeux (1985) has pointed out that genes are not likely to be responsible for the diversity and specificity of all synaptic connections. Studies of the eye in a simple organism such as the water flea reveal that the number of sensory neurons and the number of ganglion cells with which they make contact are the same for insects that are cloned and thus genetically identical; but the number of synapses between these cells and the shape of their axonal branches vary from one clone to another (Macagno, Lopresti, and Levinthal 1973). When we consider the more complicated nervous system of mammals, the variability is only likely to increase.

Considerable research has documented environmental influences on brain development. Animal experiments have shown that enriched environments increase the size of the cortex, the density of glial cells, and the density and number of synaptic connections (Diamond 1984). Research with the visual system also has shown that sensory deprivation early in life can lead to severe abnormalities in visual cortical cells and produce blindness, but deprivation later in life does not produce ill effects (Kratz, Spear, and Smith 1976).

Acknowledging the influences of both genes and the environment, several neuroscientists have suggested the theory that learning proceeds from a "selective, Darwinistic mechanism" (Changeux 1985; Dudai 1989; Edelman 1987). The key idea is that the brain spontaneously produces what Changeux has called *prerepresentations*, which correspond to the transient but discrete electrochemical activity of an assembly of neurons. These autonomous patterns of activity are endogenous; that is, they exist apart from any sensory input. Inputs from sensory receptors and the cortical maps they project to create a *primary percept*. Percepts interact with prerepresentations of the brain to produce resonant states in which the prerepresentation is stabilized and a memory or *mental image* is formed. For resonance to occur, a certain degree of "match-up" has to exist between the percept and the prerepresentation; too much dissonance will interfere with learning. Resonance is produced when neurons fire in phase with one another, creating a temporally coupled volley of activity (Von der Malsburg and Willshaw 1981). This resonant pattern reinforces itself through a strengthening of synaptic connections between neurons in the assembly, thus "storing" the memory.

Mental representations are not the direct product of sensory input, nor are they hardwired into functional circuits by the genome. Rather, stable neural states are selected from among a variety of spontaneously generated neural activation patterns through an interaction with sensory input. The idea is that prerepresentations are genetically inherited, but only those endogenous activation patterns that prove to be functionally useful through experience are stabilized to form the mental structures of the mind. This stabilization process occurs during the critical period of neural development when some neural connections are stabilized into functioning circuits and other synaptic contacts regress. Presumably neural architecture can be influenced by evolution in the same way as other physical traits. Those prerepresentations that provide the right kind of raw material to form mental representations give the animal a selective advantage.

Although these ideas are likely to remain speculative for some time, they nevertheless have directed considerable interest to the level of cell assemblies. In the study of complex representations, most neuroscientists now recognize the need to monitor the activity of large assemblies of neurons, not just record the activity of single cells. Techniques to provide such functional mapping do not yet exist, although several new experimental procedures, such as optical dyes (Grinvald et al. 1986) or mapping brain activity using multiple electrodes (Abeles and Gerstein 1988), hold promise.

7.4 *Neuropsychology*

Theoretical and Methodological Considerations

The study of the neural substrates of higher cognitive functions, such as memory, language, or consciousness, is often referred to as *neuropsychology*. Traditionally, neuropsychologists have studied how cognitive functions are physically implemented by studying people or animals with damaged brains. Performance by a person with damage to a known area in the brain can be compared with performance by people with intact brains or damage to other areas. The differences in performance can be used to generate or test theories about the roles played by various areas of the brain in cognitive functioning. In recent years this line of research has been vigorously pursued as part of the framework of cognitive science for several reasons. First, neurological patients sometimes exhibit strikingly selective cognitive deficits, which suggest that some of the detailed representations and processes that have been hypothesized at the cognitive level of analysis are implemented in localized areas of the brain. Second, recent methodological advances in cognitive psychology have made possible much more detailed assessments of the cognitive impairments of neurological patients. Third, contemporary cognitive scientific theories have generated a host of new neuropsychological theories. The flowering of neuropsychological research within cognitive science has also been accelerated by the development of new methods for observing the working brain. Localized patterns of activity during cognition can now be observed in the intact or damaged brain using methods such as magnetic resonance imaging (MRI), positron emission tomography (PET), or the recording of event-related potentials (ERP). Although their spatial and temporal resolution is limited, these methods allow researchers to observe directly some of what must be inferred in the traditional method of correlating cognitive performance with the locations of brain lesions.

Traditional and much of contemporary neuropsychology is a top-down enterprise, in which theories from the cognitive level of analysis are brought to bear on neuropsychological data. Advances in our knowledge of neural circuitry and the growth of connectionist and neural network models are beginning to bring bottom-up considerations to the field, however. Neuropsychology researchers must now pay attention to relevant data about neural circuits, and they must consider network models in addition to more classical cognitive models, which are often couched at a higher level of analysis and employ highly structured representations and rules. Currently, there is lively debate within neuropsychology over the relative weight to be given to traditional cognitive models and their supporting data versus network models and neurobiological data. Many traditional cognitive models, of course, were not designed with any general theory of neural computation in mind, and the gross anatomical data about brain lesions traditionally associated with neuropsychology say little about the organization of neural circuits. Theorists who are oriented toward the reduction of cognitive models to neural models (Churchland 1986; Sejnowski 1986) tend to argue that all neuropsychological theories should be based at least in part on well-supported characteristics of neural computation and, where possible, microstructural neuroscientific data.

The detailed study of *acquired dyslexia* (the loss of reading skills through brain damage) represents one of the more successful examples of the top-down approach to neuropsychological research (Coltheart, Patterson, and Marshall 1987; Patterson, Marshall, and Coltheart 1985). Researchers have systematically analyzed the patterns of reading errors produced by dyslexic patients to gain insight into how the reading system is structured. The overall structure of a theoretical model resulting from this line of research can typically be represented as a flowchart, such as the one in figure 7.14. Each box in the chart represents a computational module, or component, that takes one or more representations as input and maps that input onto an output representation. The label in a box describes the output representation that is computed by that module. The arrows in the chart represent the flow of information among the modules. Figure 7.14 represents a model, proposed by Marshall and his colleagues (Marshall and Newcombe 1981), which contains two parallel information-processing routes for the reading of individual words. One route is based on sight vocabulary and is known as the *lexical* or *direct* route. The other route, called the *phonological* route, relies on regularities in the correspondence between spelling and sound and on the morphological and phonological structure of spoken language. This route can "sound out" a word in the absence of a direct processing route.

In the current model word recognition begins with an analysis of visual features to provide information for identifying letters or whole words. In the phonological route letters are grouped or parsed together into syllabic units, which are then coverted to the proper phonemes. This conversion process is based on rules or regularities acquired from spelling experience (Venezky 1970). The phonological outputs of the conversion process are then passed to a further process that blends them back together to produce the whole word. The early stages of development of the phonological route are evident when young readers sound out words. In the lexical route words also are parsed in the early stages to identify the base root of a word (e.g., *antiabortion* becomes *anti/abort/ion*). Then lexical memory is accessed to recall the meaning and proper pronunciation of the base and any prefixes and suffixes. The word parts are then resynthesized for proper pronunciation. The two routes are hypothesized to operate in

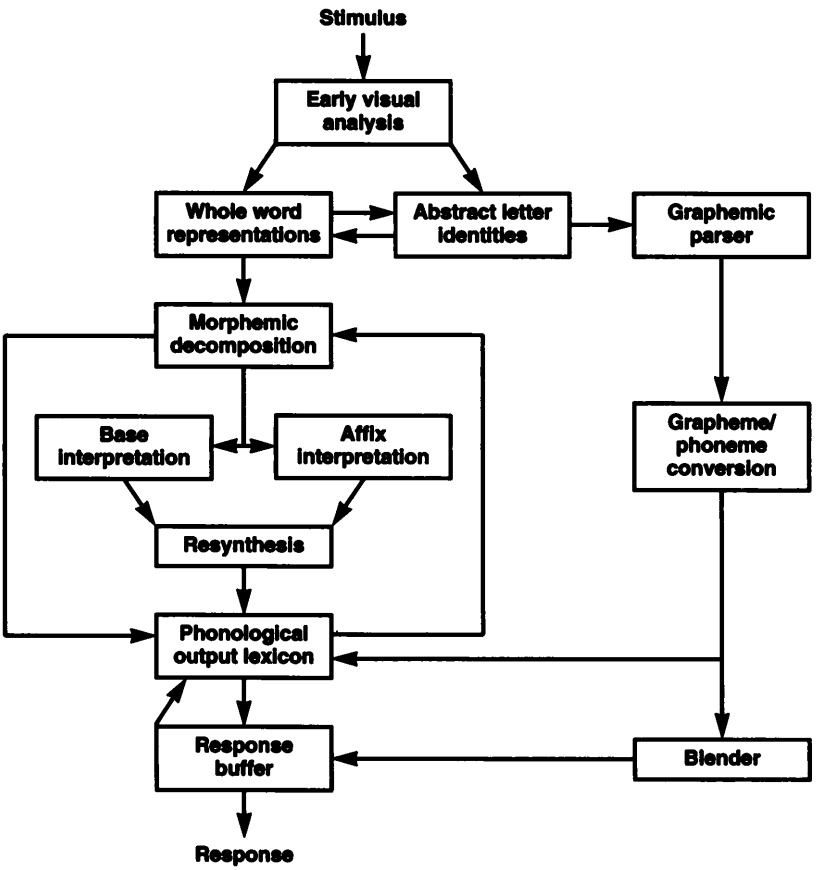

Figure 7.14
Symbolic model of reading

parallel so that whichever subsystem completes the word identification process first will control how the word is pronounced.

This model has had considerable success in accounting for the variable reading performance of acquired dyslexic patients. For example, patients with a condition called *phonological dyslexia* have difficulty pronouncing unfamiliar words, even simple nonwords such as *troat*, but otherwise may have no trouble reading. In this case the phonological route is thought to have been damaged, leaving intact the lexical route and its access to a previously learned reading vocabulary. The opposite problem is thought to exist for a condition called *surface dyslexia*. In this case the patient cannot access lexical information and so must sound out every word by reading through the phonological route. Because the phonological route follows spelling rules in deriving a pronunciation, surface dyslexics frequently mispronounce words with unique pronunciations, such as "yacht," or words that have some irregularity, such as *come* or *have* (the rule says that *e* at the end of the word should make the vowel long). Surface dyslexics also have trouble understanding what they read because access to lexical memory is impaired.

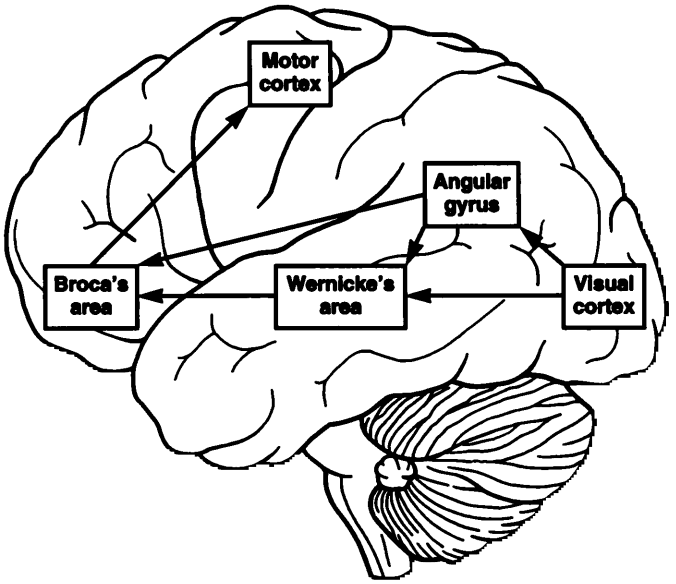

Figure 7.15
The Wernicke-Geschwind model of reading as represented in different cortical regions

Models of this kind have been mapped onto the brain (see figure 7.15) by correlating the location of brain damage and patterns of reading impairment with the processing components of the models (Benson and Geschwind 1982). Visual input is processed along two parallel pathways from occipital to parietal, temporal, and then frontal regions. In the phonological route visual information is converted to an auditory code as the word form passes through language regions in the left hemisphere. This conversion has been determined to take place in a region of the parietal lobe called the angular gyrus, since patients who suffer damage to this region lose the ability to sound out words. The lexical route is presumed to circumvent the angular gyrus and to access semantic information directly from the temporal lobe, based on the visual word form. Motor programs for controlling speech output are hypothesized to be located in a part of the frontal cortex known as Broca's area (see the discussion of language below).

Regardless of the research strategy that is followed, neuropsychology is an extraordinarily difficult enterprise. The neuropsychologist faces three sorts of methodological problems.

The first problem is shared with all cognitive scientists. Like any cognitive scientist, the neuropsychologist typically has a computational theory of the cognitive process under study. Whether this theory is a traditional symbolic theory or a more microstructural connectionist theory, it can be wrong. The proposed representations, components, or computational processes might be incorrect.

The second problem is shared with all cognitive scientists who are interested in the brain. In addition to a computational theory of some cognitive process, the neuropsychologist must have a theory of how the computational model is realized in the brain. The computational theory does not dictate the physical-realization theory.

Given the current state of our knowledge, there is almost always more than one possibility for mapping the computational theory onto the brain. A simple mapping that associates parts of the computational model with discrete parts of the brain would be incorrect, for example, if one computational component were implemented in several anatomically distinct areas of the brain rather than in one brain "center."

The cognitive neuroscientist is thus juggling two theories, each of which can fail in various ways, making it difficult to interpret data and to assign credit or blame when the data suggest a revision of the theories. Many researchers hope that this double difficulty of cognitive neuroscience will be converted into an advantage as knowledge advances. Computational and neural models, supported by behavioral and neurobiological data, should ultimately become mutually constraining. Research at one level of analysis should strongly suggest or rule out possibilities at other levels. Arguments over the merits of cognitively versus neurally inspired models aside, the prospect of converging evidence from different levels of analysis suggests that all levels should be pursued simultaneously and as cooperatively as possible in what Churchland (1986) calls a coevolutionary research strategy. For some higher cognitive processes the current problem is that not enough reliable knowledge is available at lower levels of analysis to strongly constrain the higher levels, hampering the dynamic of coevolutionary research.

The third problem for neuropsychology is specific to working with subjects who have suffered brain injury or have other kinds of neurological abnormalities. The attempt to understand the workings of an extraordinarily complex system by observing its behavior when parts of it have been capriciously damaged is obviously fraught with difficulties. To begin with, the neuropsychologist faces the additional burden of developing a theory of how the computational model of normal cognition performs when it is "lesioned" and of mapping these cognitive lesions onto physical lesions. The general empirical strategy in neuropsychology has been to show that damage to a specific brain area selectively impairs some aspect of performance and spares other aspects (which in turn may be impaired by damage to another brain area). When brain injuries *dissociate* various aspects of normal cognitive performance from each other, it becomes possible to map the injured brain areas onto the components of cognitive models that are responsible for the aspects of performance that are impaired or spared. Even in cases where the dissociation is complete, however, the mappings proposed are not logically necessary (Shallice 1988). For example, Churchland and Sejnowski (1992) argue that "lesioned" connectionist networks can exhibit dissociated performance even though they do not consist of discrete modules that are interconnected by pathways that carry different representations. One way this can happen is that, as the weights on hidden units are adjusted during learning, units can come to specialize on one or another aspect of the structure of a domain (e.g., vowels vs. consonants). If units with similar specializations are physically clustered, they could be selectively lesioned, causing a selective performance impairment.

The problems with interpreting neuropsychological data are further complicated by other factors. Naturally occurring brain damage (typically caused by stroke, trauma to the head, or infection) does not respect computational categories, and it is usually quite diffuse, involving several major areas. Localized lesions are rare. When patients with similar cognitive performance are grouped together for purposes of study, the groups are never homogeneous. The subjects usually have a variety of lesions of different sizes and in somewhat different locations. Even in a fairly homogeneous group, there might

still be considerable individual differences in the organization of cortical maps (see the earlier discussion of topographical maps). Imaging techniques used to identify brain injury also have limitations. Computerized tomography (CT), the most frequently used brain scan, cannot image a lesion smaller than 1 to 2 centimeters and does a poor job in identifying subcortical damage. Another factor to consider is the cause of the damage. Grouping patients with different types of injury is problematic. Head trauma almost always produces diffuse damage; tumors can produce electrical abnormalities; and brain function can be affected by related medical conditions, such as cynotic heart disease, which can slow down metabolic activity by depleting oxygen supply. Many patients with brain damage also take psychotropic drugs, such as the seizure preventive Tegretol, yet the effects of these drugs on cognitive functions are not well understood. Finally, the patient's age at the time of the injury and his/her time for recovery (the interval between when the injury occurred and when testing took place) can be important factors in performance. Taken together, these concerns constitute a formidable challenge to the neuropsychologist. No study of a group of brain-damaged patients ever controls for all the factors because too many patients would be eliminated from the group.

One approach to this problem has been to avoid studies that involve groups of patients and instead to make detailed case studies of individuals with brain damage. Case-history studies have attempted either to provide a comprehensive assessment of an individual's cognitive functions (the approach pioneered in the former Soviet Union by Luria) or to analyze in detail a specific cognitive function (the British approach recently championed in this country by Caramazza). There are strong opinions on both sides about the relative merits of group or case studies in neuropsychology (Caramazza and McCloskey 1988).

As research progresses, the challenges of neuropsychology should be met by a combination of improved techniques, such as the use of high-resolution brain scans and the expanding database of carefully studied subjects, and convergent data from other sources. One area where the value of converging evidence can be seen is color vision and cognition. Certain aspects of color vision performance, such as color matching, discrimination, and naming, have been very well characterized in laboratory studies. The underlying neurophysiology up to the primary visual cortex is also well understood, and it fits together very well with the perceptual data. Although the basic character of the perceptual data has been known for over one hundred years, the fit with the physiological results came with a fertile period of coevolution that began with the development of modern neurophysiological methods, such as the techniques for recording the responses of single neurons. Given that we have a good handle on some of the phenomena of color perception and cognition and that we know a good deal about how color is initially represented in the cortex, we have a good chance of being able to put neuropsychological observations about color to good theoretical use. Davidoff (1991), for example, integrates a wide range of neurophysiological and perceptual-cognitive data with findings concerning various impairments involving color, such as the selective, total loss of color vision caused by cortical injury (acquired achromatopsia) or the selective loss of the ability to name colors when other aspects of color perception and cognition are intact (color anomia). As converging sources of evidence accumulate in other areas of research, neuropsychological data will make a clearer contribution to the overall scientific picture in those areas as well.

Memory

Memory is a good illustration of the study of a cognitive faculty for which neither the cognitive nor the neurobiological theory and data are as detailed or secure as they are for color vision. The contemporary neuropsychological study of human memory began with the study of HM, a patient who underwent a radical surgical procedure to control his constant epileptic seizures. Medications had proved ineffective, and so the parts of HM's brain that were producing the seizure activity were removed (Scoville and Milner 1957). To everyone's surprise, when HM recovered from surgery, he was profoundly amnesic and could not remember events that occurred after the surgery. HM had no trouble recalling events from his childhood or utilizing information that he had learned prior to the surgery. His short-term memory skills were normal. Therefore, the surgery affected his ability to store or retrieve new long-term memories. HM has a disorder called *anterograde amnesia*, the inability to recall events that occur after the onset of the amnesia. (*Retrograde amnesia* refers to the loss of memory for events that occurred prior to the amnesia.) For HM the amnesia is so profound that he cannot remember what he was doing even a few minutes ago. He is forever stuck in the present, or as he himself described it, "Every day is alone in itself, whatever enjoyment I've had, and whatever sorrow I've had.... It's like waking from a dream. I just don't remember" (Milner, Corkin, and Teuber 1968, 216).

The case of HM redirected attention to the study of storage and retrieval processes as well as different aspects of long-term memory. As we discussed in chapter 3, HM and other amnesic patients like him have difficulty with semantic and episodic memory. Semantic memory refers to memory for facts, such as the meaning of a word or the birth date of a grandson, and episodic memory refers to recall of specific autobiographical events, such as a boat trip last Fourth of July or yesterday's breakfast (Squire 1987). In our sketch of the cognitive architecture we classified both of these kinds of memory as declarative, the kind of memory that can be learned in just one trial or experience, and can be voluntarily and consciously accessed, usually through more than one sensory modality. We contrasted declarative memory with procedural memories, which are built up over repeated exposures, are manifested by changes in performance rather than by conscious recollection, and are modality specific. Figure 7.16 presents a classification scheme for types of memory.

As we pointed out in chapter 3, studies suggest that the procedural memory system is intact for amnesic patients. Milner (1965) was the first to demonstrate that HM could

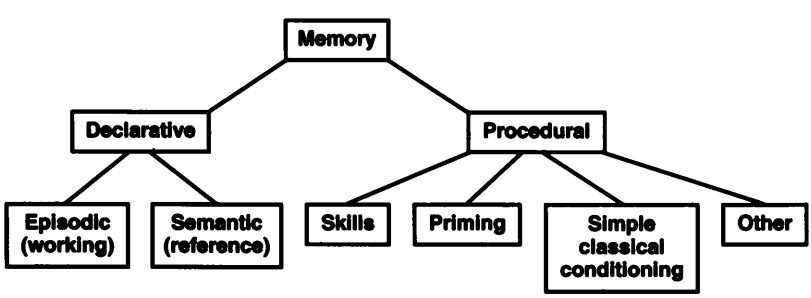

Figure 7.16
Model of memory systems. (From Squire 1987.)

learn new skills by demonstrating his improved performance on a mirror-tracing task after daily practice even though each time he practiced, he had no recollection that he had done the task before. Since this first study, improved learning has been demonstrated for many other behavioral domains, including maze learning, hand-eye tracking, mirror reading, the Tower of Hanoi problem, and jigsaw puzzles (Squire 1987). One study also reported that amnesic patients could be physiologically conditioned using a paradigm in which they learned to blink their eyes when they heard a tone (Weiskrantz and Warrington 1979).

HM and other amnesic patients also show repetition priming effects on tasks such as the word-stem completion test described in chapter 3. In this task the subject is first exposed to a list of words, such as MOTEL or ABSENCE. Later the subject sees a list of word fragments, such as MOT__, ABS__, or CHE__, and is asked to say the first word that comes to mind that starts with the same letters. At least ten different words could be constructed from each fragment, and so by chance subjects could guess the primed word at most 10 percent of the time. Both amnesics and controls perform well above the guessing rate and show similar priming effects by responding with words from the priming list about 60 to 70 percent of the time, even if they are discouraged from memorizing the first list (Graf, Squire, and Mandler 1984). When the task is converted into a standard declarative memory task by instructing subjects to use the word fragments as cues to the recently presented words, amnesic subjects are unable to make use of the cues. Such priming is temporary and disappears if a delay, ranging from several hours to days, is imposed before showing the word fragments. Schacter's hypothesis that priming affects perceptual representation systems that are not part of the declarative memory system was discussed in chapter 3.

The marked difference in amnesic performance between declarative and procedural memory tasks provides strong support for considering these two types of knowledge to be organized into different neurobiological systems. Squire (1987) has described procedural memory as the mental history of accumulated experiences "tied to and expressible through activation of the particular processing structures engaged by learning tasks. It is acquired and retained by virtue of the plasticity inherent in these structures" (p. 162). Such learning is hypothesized to be contained in the specific neural circuits that are used to complete the procedural task. Through repeated use, structural (e.g., dendritic spines) and neurochemical (e.g., calcium ion channels) changes occur in the neurons that facilitate the synaptic connections, producing a functional neural circuit capable of computing the desired behavioral response. This learning mechanism is hypothesized to be present in neural circuits throughout the CNS. The theory that procedural learning involves local changes in the efficiency of neural circuits may explain why procedural memory tends to be highly specific. For example, procedural memories typically do not transfer well across sensory modalities. A simple case is the priming paradigm just described. The priming effect is reduced when the initial presentation of the prime words is auditory rather than visual. Declarative memories, however, are organized, stored, and recalled together from many different sensory modalities; for example, memories of a Fourth of July from childhood might be triggered by the smell of cotton candy or by the sound of a familiar song from that time. To consolidate the memory of a single event appears to require a central structure that functions to associate information from different sensory modalities. Hippocampal and amygdaloid structures are thought to serve this function. This consolidation process

must be contained in neuroanatomical structures that are physically separate from those where the long-term memories are stored, since HM could recall declarative information that was acquired before the surgery.

Neural Model of a Memory System Several different brain structures have been implicated in memory functioning. HM's hippocampus from both hemispheres was surgically removed, along with both amygdalae and parts of both temporal lobes. Consequently, intense interest has focused on the hippocampus and amygdala as the site for long-term declarative memory consolidation (Amaral 1987). However, other cortical regions also have been implicated in amnesia. Long-term alcoholism accompanied by thiamine deficiencies can produce a neurological syndrome called Korsakoff's disease. Korsakoff's patients suffer from damage to the thalamus (dorsomedial and anterior nuclei) and occasionally to the mammillary bodies of the hypothalamus, although researchers now believe the thalamic damage is most likely to cause the memory loss (Zola-Morgan and Squire 1985).

Mishkin and his colleagues have proposed a model in which the medial temporal (hippocampus and amygdala) and diencephalic (mammillary bodies and thalamus) regions function together as a memory system (Mishkin and Appenzeller 1987). Mishkin's model incorporates information not only about these structures but also about brain regions that are connected to them, such as the prefrontal cortex, basal forebrain, and temporal and parietal structures (see figure 7.17). In his model each region serves different roles in the memory process. Sensory information moves from the primary sensory cortical regions to the associative cortex, where long-term memories are stabilized by concurrent feedback activity from the hippocampus and amygdala. This feedback circuit has two loops: an indirect route that passes from diencephalic structures through the prefrontal cortex and then the basal forebrain and another route that connects more directly through the basal forebrain. The prefrontal cortex has been hypothesized to organize and associate behavioral responses based on the current sensory input. The hippocampus and amygdala are thought to function as a working memory system, to help consolidate new perceptions from short- to long-term memory and also to assist in memory recall by associating the many features of an episodic memory across different sensory modalities. Damage to any part of the system will disrupt functioning, but not to the same degree. These ideas about the functional role of each region still remain tentative, but more systematic study involving a broader sample of standardized behavioral tasks may help to sort these questions out.

Language

We have already been exploring the idea that different parts of the brain are specialized for particular cognitive functions. In 1861 two French physicians, Auburtin (1825–1893) and Broca (1824–1880), described a patient who had suffered speech loss, a disorder also called an *aphasia*, as the result of damage to part of the frontal lobe (see figure 7.18). Broca eventually examined a total of twenty-three individuals who had language disorders and concluded that in each case the frontal lobe of the left hemisphere had been damaged. Further, he noticed that only speech production was impaired; comprehension remained intact. Today this region is called *Broca's area*, and patients who suffer speech loss from damage to it have *Broca's aphasia*.

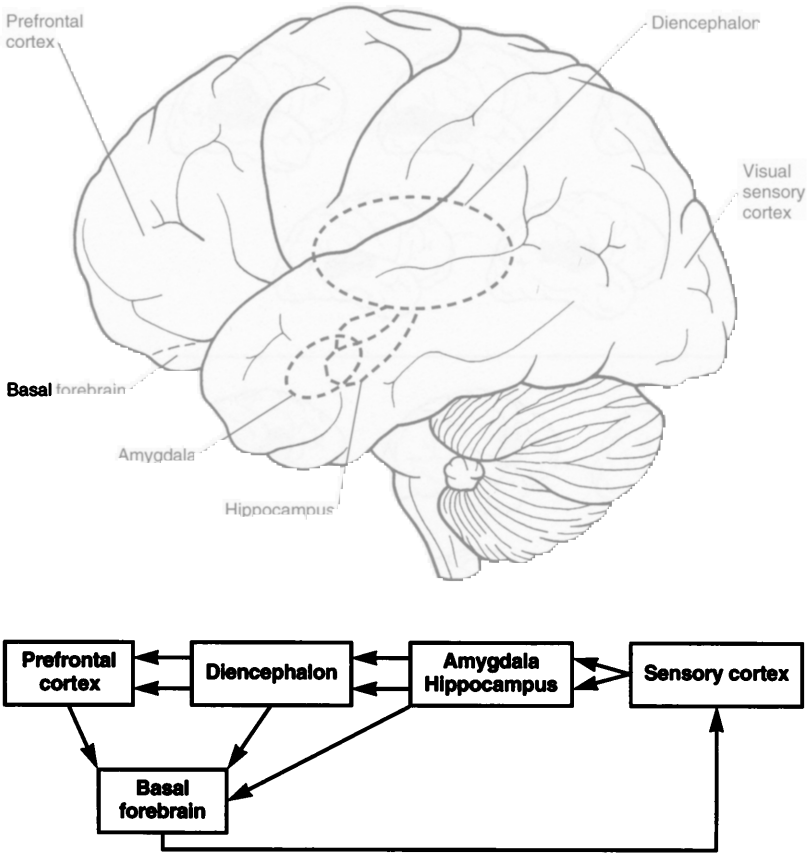

Figure 7.17
The flowchart and anatomical locations of a model for visual memory functions

Broca's aphasia patients are characterized by slow and effortful speech in which each word can require several seconds to produce and speech sounds are slurred or rearranged. For example, Gardner (1974) reports this conversation with a Broca's aphasia patient:

> I asked Mr. Ford about his work before he entered the hospital.
> "I'm a sig ... no ... man ... uh, well, ... again."
> "Let me help you," I interjected. "You were a signal ..."
> "A sign-nal man ... right," Ford completed my phrase triumphantly.
> ... "I see. Could you tell me, Mr. Ford, what you've been doing in the hospital?"
> "Yes, sure. Me go, er, uh, P. T. nine o'cot, speech ... two times ... read ... wr ... ripe, er, rike, er, write ... practice ... get-ting better."
> "And have you been going home on weekends?"
> "Why, yes ... Thursday, er, er, er, no, er, Friday ... Bar-ba-ra ... wife ... and, oh, car ... drive ... purnpike ... you know ... rest and ... tee-vee." (Gardner 1974, 60–61)

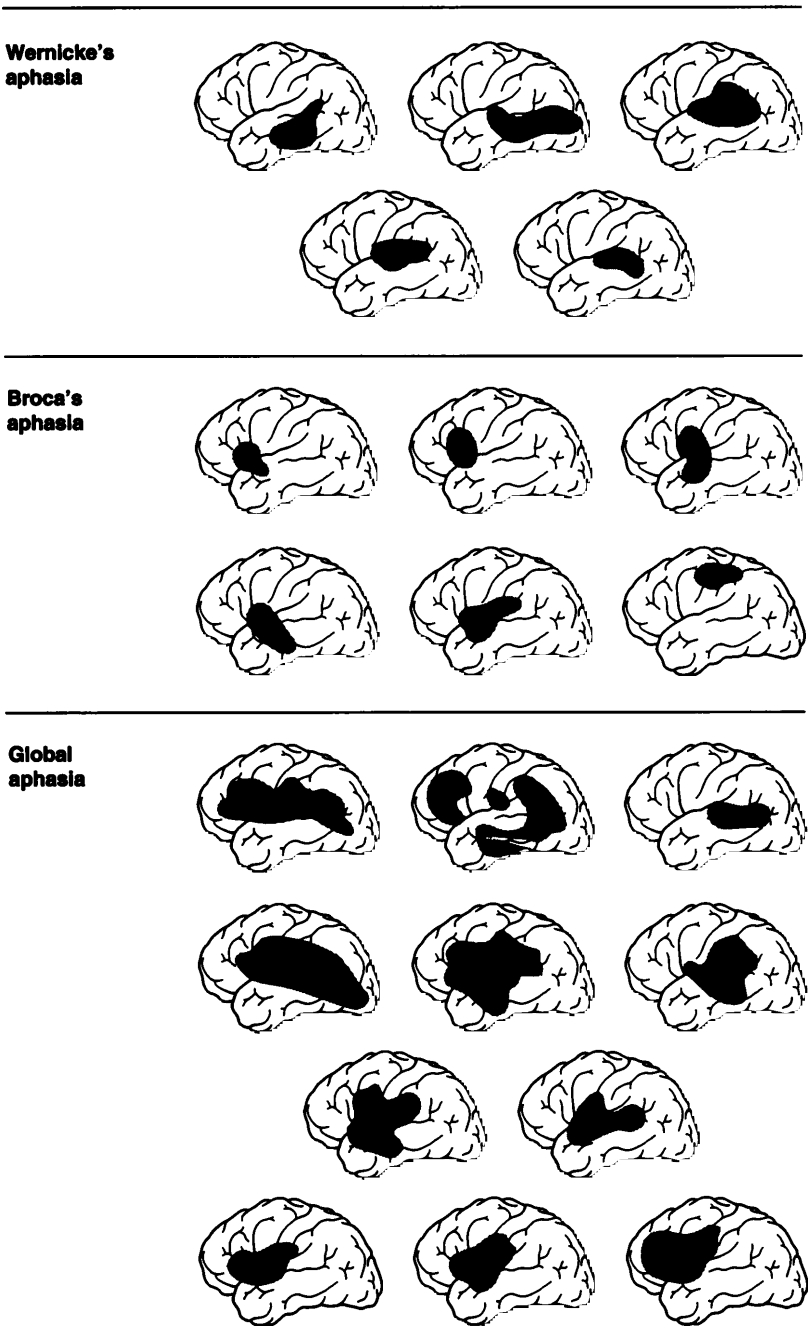

Figure 7.18
Figures showing the site of brain damage (dark areas) for different individuals. Top group suffered from Wernicke's aphasia, middle group from Broca's aphasia, and bottom group from Global aphasia. (From Kolb and Whishaw 1990.)

About ten years after Broca's report, Wernicke described a second language center located in the temporal lobe. Whereas patients with Broca's aphasia talk in a slow, deliberate manner, using simple grammatical structures, those with Wernicke's aphasia have poor speech comprehension, and although speech is fluent, they often confuse the sounds of words and mix up syllables to create neologisms (new words) or scramble phrases together to produce word salad. Again, another example from Gardner's book:

> "What brings you to the hospital?", I asked the 72-year-old retired butcher four weeks after his admission to the hospital.
> "Boy, I'm sweating, I'm nervous, you know, once in a while I get caught up, I can't mention the tarripoi, a month ago, quite a little, I've done a lot well, I impose a lot, while, one the other hand, you know what I mean, I have to run around, look it over, trebbin and all that sort of stuff." (Gardner 1974, 68)

In addition, the writing skills of Wernicke's aphasics are usually impaired.

Since Wernicke's time, neurologists and neuropsychologists have continued a functionalist approach to the study of neurolinguistics. The method involves a careful examination of the linguistic skills of brain-damaged patients to correlate their speech impairments with the site of their brain trauma. Goodglass and Kaplan (1972) have produced a classification system that is commonly used to analyze linguistic functions. This system includes two types of comprehension disorders, visual and auditory, and eight types of expressive disorders, which cover areas such as articulation, grammar, fluency, and writing. Most aphasic patients have impairment in many of these categories, involving both expressive and receptive functions. Although many neuropsychologists who study aphasia hope and expect that someday each of these functions will be localized to discrete anatomical locations, so far this goal has not been reached. One reason may be the extent of most patients' brain damage. Aphasia is commonly caused by strokes. The middle cerebral artery, a blood vessel in the brain that nourishes these language regions, is particularly susceptible to arteriosclerosis, a disease in which blood vessels thicken and are weakened with age. When a blood vessel is occluded or bursts, large portions of brain tissue die, typically producing a lesion that is not localized to a specific functional area.

Other explanations for the difficulty in localizing such functions may be either that the brain does not organize linguistic processes by Goodglass and Kaplan's categories (Marshall 1986) or that such functions are more broadly distributed and not localized. In practice only about 60 percent of aphasic patients exhibit patterns of linguistic impairment that fit into the current classification schemes.

Linguistic aphasiology research has adopted an information-processing approach to provide a much more detailed description of linguistic function (Caplan 1987). Using theoretical concepts derived from linguistics and cognitive psychology, aphasiologists have studied the impaired linguistic performance of aphasic patients to map their deficits into subcomponents of a language-processing system. Many interesting refinements of linguistic theory have resulted from this work. For example, several studies have examined the role syntactic structures play in sentence comprehension. Consider these three sentences:

(1) The girl is chasing the big dog.

(2) The dog the girl is chasing is big.

(3) The dog the girl is chasing is barking.

Although sentences (1) and (2) convey the same meaning, the order of the nouns is reversed in sentence (2), requiring the reader to use syntactic information to assign the proper thematic roles to the dog and the girl to determine who is doing the chasing. Syntactic information also is useful in sentence (3), although lexical-pragmatic information can be used to determine thematic roles since girls do not bark. When sentences like these are given to Broca's aphasia patients, they have a very difficult time comprehending sentence (2) but not (1) or (3), suggesting the presence of several different systems involved in language comprehension (Caramazza and Zurif 1976). However, Broca's aphasics are not completely impaired in their grammatical judgments. Linebarger, Schwartz, and Saffran (1983) reported that Broca's aphasics who were severely impaired in their syntactic comprehension were still able to make reasonably good judgments about whether sentences were grammatically correct. This finding suggests that grammatical systems must be broken down further into parsing operations that determine constituency and a separate stage in which the sentence structure created by the parser is interpreted.

The issue of localizing linguistic functions to particular brain regions is problematic. Although considerable evidence suggests the principle of gross localization of major linguistic functions, more elementary stages of linguistic processing appear to be distributed and to vary between individuals. Some of the most problematic evidence comes from studies conducted by Ojemann and his colleagues (Ojemann 1983), who attempted to map speech zones during neurosurgery. Ojemann used a weak electrical current to stimulate selected surface areas of the cortex while patients were engaged in verbal tasks. During the procedure patients remained awake and alert because there are no pain sensors on the surface of the brain. If their performance was disrupted, then the site of the electrode indicated a speech zone. Their results supported the general concept of Broca's and Wernicke's language zones; however, stimulation in both regions had quite similar effects, disrupting both expressive and receptive functions. In addition, the boundaries varied considerably from one individual to another, and many other regions outside of these zones also affected linguistic functions. Furthermore, many different tasks, such as phoneme perception (distinguishing *ba* from *da*) or copying orofacial movements (such as sticking the tongue out), could be disrupted from the same site.

Other brain structures besides the cortex have often been overlooked when discussing language functions. In 1866 Hughlings-Jackson (1932) was the first to suggest that subcortical structures are important to normal language function. Recent studies have demonstrated language disorders in patients with thalamic lesions (Ojemann 1975) and in Parkinson's patients who have basal ganglia damage (Lieberman, Friedman, and Feldman 1990). Together these data suggest that the language system is extremely complex, involving many cortical and subcortical structures.

Left Brain/Right Brain—Cerebral Dominance

A curious feature of the brain's organization is that the symmetrical halves of the brain respond to sensory input from the opposite or contralateral side of the body. The somatosensory system is almost completely crossed so that the right half of the brain controls the left side of the body and vice versa. The visual system is slightly more complicated (see figure 7.19), with the nasal (near the nose) portions of the visual field crossing while the peripheral visual fields remain ipsilateral (same sided). Thus, if the eyes are focused straight ahead, the left visual field is mapped onto the right visual

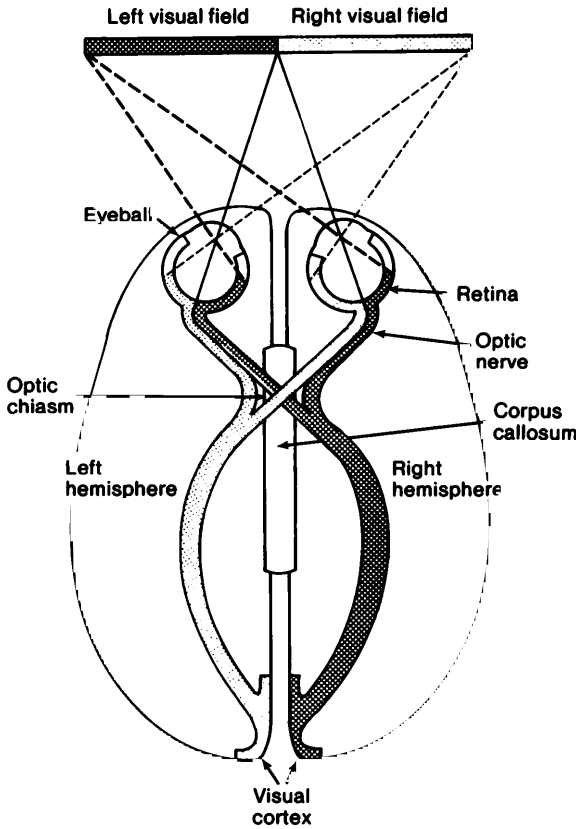

Figure 7.19
Central visual pathways. Light from the left visual field projects to the left eye's nasal and right eye's temporal parts of the retina. These regions send optic fibers to the right cortical hemisphere. Similarly light from the right visual field projects to the left cortical hemisphere.

cortex and vice versa. The auditory system also is contralaterally organized, although ipsilateral projections exist as well.

Another unusual feature is that one cerebral hemisphere is usually dominant for particular behavioral functions. For example, the vast majority of people show a right-handed preference, suggesting left hemisphere cerebral dominance for manual motor functions. Many other cognitive functions, such as language, are also lateralized to the left hemisphere, although some, such as visual-spatial abilities, show a right hemisphere preference (Bryden 1982). Why higher cognitive functions, such as language, should lateralize to a dominant hemisphere has been the subject of intense speculation and recent biological investigation (Geschwind and Galaburda 1984). One hypothesis suggests that the two hemispheres are not as physically symmetrical as once thought and that important anatomical and cellular differences may be at the heart of certain functional differences. For example, Geschwind and Levitsky (1968) have shown that a region known as the *planum temporale*, located in the temporal lobe near the primary auditory corex (see figure 7.20), is larger in the left hemisphere for about 80 percent of the brains examined, suggesting a possible role in the left hemisphere's dominance for linguistic functions.

Split-Brain Studies Since the 1940s *commissurotomy* has been used as a last resort treatment for some patients with intractable seizures. This is a rare neurosurgical procedure in which the fibers of the corpus callosum that connect the two cerebral hemispheres of the brain are cut. The surgery is thought to prevent the spreading of epileptic discharge from one hemisphere to the other, thus helping to control seizure activity. Two groups of patients who have had commissurotomies have been carefully studied, revealing important differences in the way the two hemispheres function. The California group consists of about two dozen patients who underwent complete commissurotomies in the 1960s. They were extensively examined by Sperry and his colleagues (Gazzaniga 1970; Sperry 1968). During the 1970s a second group of patients underwent partial commissurotomies at Dartmouth Medical School and were studied by Gazzaniga and his colleagues (Gazzaniga and LeDoux 1978).

Study procedures took advantage of the fact that sensory systems are contralaterally organized for the two hemispheres. If visual information was only briefly displayed to the left visual field (so as to avoid eye movements that would shift the field of vision), or if objects were handled only by the left hand, the sensory information would be conducted only to the right hemisphere. With the corpus callosum severed, information from one hemisphere could not be transferred cortically to the other hemisphere, providing the opportunity to directly examine the cognitive processing capabilities of each hemisphere.

In general, split-brain studies confirmed that the control of speech is localized to the left hemisphere (LH) for most people. For most subjects the LH was able to respond verbally to questions or other visual displays, whereas the right hemisphere (RH) was unable to do so. Gazzaniga (1983) reported that only three out of twenty-eight patients from the Dartmouth group showed any evidence of RH linguistic function, and in all three cases evidence suggested the RH may have assumed a role in language because of LH damage early in life. According to Gazzaniga, only two patients in the California group showed evidence of RH language function, although Zaidel disagreed and claimed that as many as six patients had RH language (Zaidel 1983). Part of the disagreement may be due to differences in how linguistic responses are defined. None of the patients from the California group were able to write or speak with the RH. However, several were able to complete some linguistic comprehension tasks with the RH, such as matching simple, concrete nouns spoken out loud with pictures selected by the left hand or following oral commands to make a fist or raise the left arm. Zaidel also devised a contact lens that would block one part of the visual field so that patients could visually examine objects for a longer period of time and still ensure that the information was selectively processed only by one hemisphere. Using this lens, Zaidel extensively studied two patients and reported a surprising degree of linguistic comprehension in the RH, roughly equivalent to a ten-year-old level, although comprehension of complex sentences, such as "Before touching the red circle, pick up the green square," was still compromised. More recently Gordon (1980) reported that if the LH was occupied doing a task, the RH could respond to a verbal command, suggesting that the LH may dominate and inhibit the RH from functioning linguistically (Smith 1966).

The split-brain studies also supported the notion of the RH as specialized for visual-spatial functions, although this conclusion is based on only a small selection of RH responses. For most of the split-brain patients, the RH rarely responded at all (Churchland 1986). When there were responses, the RH was superior to the LH on such tasks

as drawing figures, analyzing part-whole relationships for geometric shapes, and manipulating spatial relationships.

Remarkably, none of these hemispheric disconnection effects are seen in patients who were born without a callosum. In very rare cases the callosum fails to form during fetal development, a condition known as callosal agenesis. The few studies conducted with these patients have shown both hemispheres able to process verbal and spatial tasks (Ettlinger et al. 1972; Saul and Sperry 1968). Although most split-brain patients cannot respond verbally to pictures shown to the RH, even though their left hand can select the correct objects, callosal agenesis patients have no such difficulty. The nervous systems of the callosal agenesis patients may have compensated during development for the lack of a corpus callosum, recruiting subcortical connections to communicate between hemispheres.

Hand Preference About 90 percent of people prefer to use their right hand when they write, eat, or need to use one hand to perform other skilled activities. Because some people are comfortable using more than one hand for different unimanual tasks, researchers now treat hand preference as a continuous variable, with some people strongly lateralized to one side and others more nearly ambidextrous. To measure an individual's hand preference, questionnaires have been constructed to sample hand use on a variety of tasks. Such tests produce scores that represent an average performance, called a laterality quotient. When Oldfield (1971) surveyed 1,000 undergraduates at the University of Edinburgh, he found that students who preferred to use their right hand did so quite consistently, whereas those who preferred the left hand tended to be more ambidextrous. This finding suggested that right-handedness is the norm, and those who are not right-handed vary to differing degrees from that normal condition.

Several neurobiological explanations of handedness have been proposed, but none has received widespread acceptance. Recent studies have demonstrated several anatomical asymmetries associated with hand preference (Kolb and Whishaw 1990). On measures of cerebral blood flow, and relative sizes of left-right cortical regions, such as the width of frontal or occipital lobes, up to 60 percent of the right-handers generally show strong lateralization, whereas left-handers show the opposite asymmetry or no differences. The functional significance of these facts, however, is hard to explain. Behavioral studies of left- and right-handers suggest that left-handers have more bilateral representation of function (Springer and Deutsch 1989). For example, a procedure known as the Wada test is often used with patients who are about to undergo neurosurgery to determine which half of the cortex controls speech production. In the Wada procedure one hemisphere is temporarily anesthetized by an injection of sodium amobarbital into the carotid artery that provides the blood supply to that hemisphere. Studies of hundreds of patients have shown that 96 percent of right-handers have speech lateralized to the left hemisphere and 4 percent to the right. However, only 70 percent of left-handers have speech lateralized to the left hemisphere; of the others, 15 percent have speech lateralized to the right and 15 percent have bilateral speech representation (Rasmussen and Milner 1977). One might then conclude that the structural asymmetries are behind these functional differences, but a careful examination of individuals suggests that is not case. At present no one has produced a convincing explanation to account for the significance of the structural differences.

A new hypothesis was proposed by Geschwind and Galaburda (1987), suggesting a hormonal theory to account for the association among left-handedness, immune

disorders, and learning disabilities. In one study, 500 strongly lateralized left-handers were compared to 900 strongly lateralized right-handers by questionnaire. Compared to right-handers, results showed that left-handers had about two and a half times the incidence of migraine headaches, thyroid conditions, and immune disorders particularly involving the gastrointestinal tract, such as celiac disease, Crohn's disease, or ulcerative colitis. Left-handers also had ten times the rate of developmental disorders, including dyslexia, stuttering, language disorders, hyperactivity, autism, and Tourette's syndrome. A second study with another 1,400 subjects replicated these results. Noticing that these disorders occur more frequently in men and that the male hormone, testosterone, is known to affect the developing nervous system in nonhuman species, Geschwind and Galaburda proposed that exposure to higher concentrations of testosterone slows the growth of the left hemisphere and affects the immune system during fetal development, producing immune malfunctions and disrupting cortical organization to affect handedness and cognitive development.

For example, microscopic studies of the brains of a small sample of dyslexics have shown a cortical symmetry between the left and right regions of the planum temporale (see figure 7.20), an area thought to be involved in language processing, whereas in normal readers this region is larger in the left hemisphere 76 to 84 percent of the time (Galaburda, Rosen, and Sherman 1989). In dyslexics the planum also has high concentrations of ectopias and dysplasias, abnormalities produced by the neural migration errors described above. Geschwind and Galaburda's hormone hypothesis suggests that the dyslexic left hemisphere is smaller and contains structural abnormalities because of exposure to high concentrations of testosterone during fetal development.

A single mechanism that could account for so many different neural conditions is appealing; however, two recent pieces of evidence have weakened this hypothesis. First, studies of neural development in animals have shown that cortical asymmetry

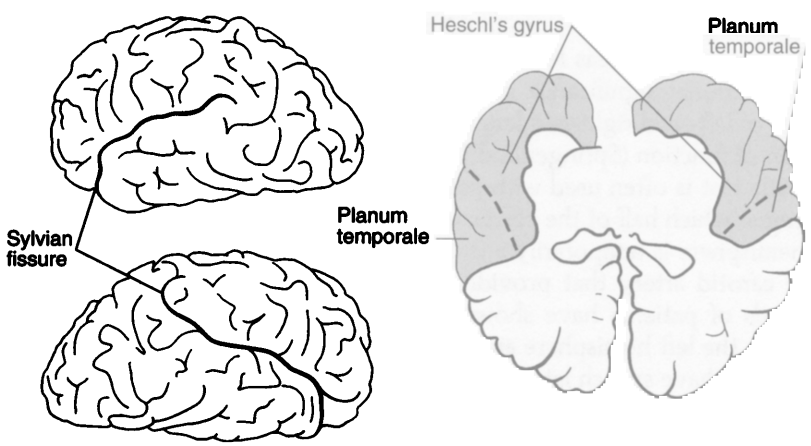

Figure 7.20
Differences in the anatomy of the two hemispheres can be found in the temporal lobe. As seen in the left figure, the Sylvian fissure rises more steeply in the right hemisphere. This difference is related to the size of the planum temporale, which is located along the surface of the Sylvian fissure (exposed in the right figure). The planum is larger in the left hemisphere. (From Kolb and Whishaw, 1990.)

develops as a result of synaptic reduction and cell death in the smaller hemisphere (Sherman, Rosen, and Galaburda 1989), suggesting that the hypothesis that testosterone inhibits left hemisphere development to produce abnormally symmetrical hemispheres is probably wrong. Apparently the normal hemispheric asymmetry is produced by neural reduction in the right hemisphere rather than greater cell proliferation in the left, and so symmetry is the result of more cells and synaptic connections remaining intact in the right hemisphere. It remains to be seen whether these results also hold true for humans, since neural cell death during human development may not be as extensive as animal studies would suggest (Huttenlocher 1990). Second, the higher incidence of some developmental disorders in males has recently been questioned. For example, several researchers (DeFries 1990; Finucci and Childs 1981) have suggested that boys are overrepresented in the dyslexic population because of sampling bias. If gender differences in some of these disorders turn out to be culturally defined and not biologically based, then the testosterone hypothesis would have to be modified or abandoned.

7.5 *Computational Neuroscience*

Computational neuroscience is a new term used for the study of the biological implementation of information processes. This research approach is an attempt to relate psychological models of behavior to neurobiological functions. Constructing computer models, which can mathematically define and simulate the computational components of neural systems, is an important part of this approach and a necessary adjunct to neurobiological experimentation. Sejnowski, Koch, and Churchland (1988) cite three reasons. First, a computer model can capture complex interactions among neural components, making the analysis of complex, nonlinear brain functions more accessible. Second, model simulations can provide novel explanations for complex behaviors. Simulations often suggest new experiments that could then test the model's predictions. Third, computer models can provide a vehicle for conducting simulated experiments that cannot be performed in vivo.

At a recent symposium on computational neuroscience (Schwartz 1990), many questions were raised about the enterprise: Can research on simple invertebrate nervous systems be generalized to more complex mammalian brains? Can the neural basis of cognition best be understood by a study of the logical functions of the brain, or must computational models also be designed with the neural "hardware" in mind? What spatial and temporal scales should be used in computational models—should modeling work focus at the synaptic-neuronal level, or should cell assemblies and cortical mapping schemes be studied? All points of view are represented in the neuroscientific community, and it is premature to predict which efforts will be most successful. However, some general guiding principles have emerged. First, known neurobiological principles are the standard against which model simulations must measure their success. Otherwise, neuroscientists are engineering designs of their own, as computer scientists would design an intelligent machine, not discovering how the brain functions; computational neuroscience is the study of "natural" intelligence, not artificial intelligence. Second, models can be realistic and incorporate as much biological detail as possible or can be simplified, providing a conceptual framework for study of algorithmic solutions. Both approaches have merit.

Connectionist Models

At present classical symbolic models cannot be reduced to neurobiological terms. We saw in chapters 2 and 3 that many connectionist models also have been developed to account for cognitive phenomena without any detailed assumptions about the underlying neural architecture. Smolensky (1988) and others have argued that connectionist models concern a subsymbolic, or microstructural, level of analysis that is lower than that of classical symbolic models but still above and in many respects autonomous from the physiological level of analysis. Nevertheless, connectionists have deliberately attempted to incorporate what is sometimes called *brain-style processing* into their models. Connectionist models are sometimes called *artificial neural networks*, suggesting that the processing units are analogous to neurons or cell assemblies in the brain. Connectionist models have in fact been rapidly assimilated into the neuroscience community (Gluck and Rumelhart 1990; Grossberg 1988; Hanson and Olson 1990; Nadel et al. 1989; Schwartz 1990).

Some neuroscientists (Crick, 1989; Shepherd, 1990a) have argued that current connectionist designs make many unrealistic assumptions, and they have encouraged network modelers to develop systems that incorporate more features of real neural architecture. On the one hand, some connectionist models have characteristics for which there is no neuroscientific evidence. The backpropagation learning algorithm, for example, requires that each forward path through a network be paired with a backward path along which error signals can be sent. In the model it is easy to imagine sending the error signals back along the very lines that carry activation forward through the net, but nerve fibers are not bidirectional, and there is currently no evidence for the precisely matched forward and backward paths that the model requires (Crick 1989; Zipser and Rumelhart 1990). On the other hand, there is ample evidence for the computational relevance of many features of real neural networks that are often not incorporated into model networks. For example, real neurons generate streams of action potentials, which have both a frequency and a phase, whereas the units in many network models generate simple numerical values, which lack phase information. Real neural networks contain multiple neurotransmitters and many distinct morphological types of neurons. Most artificial networks do not incorporate these distinctions. The net input to a unit in a typical connectionist model is a weighted sum of its inputs. In real neural networks there is evidence for microcircuits within dendritic trees that compute complex, nonlinear functions rather than simple summations. This finding supports an approach that takes the synapse, rather than the neuron, to be the basic unit of neural computation (Shepherd 1990a). Currently, network models vary widely in neural realism. The connectionist models discussed in this book fall at the abstract end of the spectrum.

In spite of attempts to make them more neurobiologically realistic, models always fail to incorporate some relevant details of actual neural networks and rest on some assumptions that are not fully supported by evidence. Limitations on neural realism are partly a matter of the need for better empirical information and partly a matter of the difficulties of simulating models on computers. A more important factor, however, is the belief that in some cases a simple model will yield greater scientific insight than a more complex one because it shows clearly how some features of the brain contribute to a particular function. In the quest to understand the brain at the physiological level of analysis, the key issue concerning any particular type of model is whether it succeeds in giving us insights into neural computation.

Abstract connectionist models are in fact yielding insights into brain function. Zipser (1990), for example, trained a connectionist network to compute the correct spatial direction of an object using the pattern of activation on the retina and the position of the eyes as inputs. Visual systems must compute this function because an object at a particular position in space casts an image onto different parts of the retina depending on the direction of gaze. The true direction of the object must be recovered by combining information about retinal position with information about eye position. The backpropagation algorithm was used to train a network consisting of input units that encoded retinal and eye position, a layer of hidden units, and output units that encoded true spatial direction. The input representations were designed on the basis of evidence about the actual inputs to the visual cortex. Following training, the response characteristics of the hidden units in the model closely resembled the electrophysiological behavior of parietal neurons recorded in macaque monkeys. These neurons were suspected of being involved in computing visual location (Zipser and Andersen 1988). This result demonstrated that the position computation can be acquired through experience, and it showed how neurons with particular response properties might arise during learning. Issues concerning the biological plausibility of backpropagation are less important than they might initially appear in this research because any learning algorithm that minimizes error in weight space (see chapter 2) will yield the same result. Although the model is extremely simple and unrealistic in some respects, its very simplicity makes for a stronger scientific result. Zipser was able to determine exactly which characteristics of the model produced hidden units with the desired response characteristics. Several other researchers have reported successes using similar approaches (Anastasio and Robinson 1989; Lehky and Sejnowski 1988).

A potentially powerful feature of this research is the use of learning algorithms to develop hypotheses about internal neural representations. Experimental techniques for recording and interpreting the patterns of activation among large sets of neurons are only just beginning to be developed and are still a long way from providing detailed information about representational states. It is difficult, therefore, to construct a representational theory directly from neurobiological data. If data is available on input and output representations, however, a model with hidden units can be simulated, and the representations developed by the hidden units can be studied to develop hypotheses about what to look for in the brain.

The vision of cooperative interaction among levels of research may be coming closer to realization. Empirical results in neurobiology can guide the construction of network models, and the results of model simulations can be used to guide further research on the brain.

Suggested Readings

The Brain (Thompson 1994) provides a very readable introduction to neuroscience. There are many good sources of information about basic neuroanatomy and neurophysiology: *Fundamental Neuroanatomy* (Nauta and Feirtag 1986) and *Synaptic Organization of the Brain* (Shepherd 1990b) are among the best. In *The Computational Brain* Churchland and Sejnowski (1992) provide a connectionist approach to computational neuroscience that includes detailed discussions of many neuroscientific details. *The Man Who Mistook His Wife for a Hat and Other Clinical Tales* (Sacks 1985) is a delightful introduction to clinical neuropsychology. *Fundamentals of Human Neuropsychology*

(Kolb and Whishaw 1990) is an excellent reference source. For a readable introduction to connectionist modeling, see *Connectionism and the Mind* (Bechtel and Abrahamsen 1991); for a more advanced review, see *Parallel Distributed Processing*, Vols. 1 and 2 (Rumelhart and McClelland 1986; McClelland and Rumelhart 1986).

References

Abeles, M., and G. L. Gerstein (1988). Detecting spatiotemporal firing patterns among simultaneously recorded single neurons. *Journal of Neurophysiology* 60, 909–924.

Amaral, D. G. (1987). Memory: The anatomical organization of candidate brain regions. In F. Plum, ed., *Handbook of physiology: Higher functions of the nervous system*. Bethesda, Md.: American Physiological Society.

Anastasio, T. J., and D. A. Robinson (1989). Distributed parallel processing in the vestibulo-oculomotor system. *Neural Computation* 1, 230–241.

Andersen, R. A., R. M. Siegel, and G. K. Essick (1987). Neurons of area 7 activated by both visual stimuli and oculomotor behavior. *Experimental Brain Research* 67, 316–322.

Aram, D. M., B. L. Ekelman, and L. L. Gillespie (1990). Reading and lateralized brain lesions in children. In K. von Euler, I. Lumberg, and G. Lennerstrand, eds., *Brain and Reading*. Hampshire, England: Macmillan.

Assaf, A. A. (1982). The sensitive period: Transfer of fixation after occlusion for strabismic amblyopia. *British Journal of Ophthalmology* 66, 64–70.

Baldi, P., and W. Heiligenberg (1988). How sensory maps could enhance resolution through ordered arrangements of broadly tuned receivers. *Biological Cybernetics* 59, 313–318.

Barlow, H. B. (1972). Single units and sensation: A neuron doctrine for perceptual psychology? *Perception* 1, 371–394.

Bastiani, M. J., C. Q. Doe, S. L. Helfand, and C. S. Goodman (1985). Neuronal specificity and growth cone guidance in grasshopper and drosophila embryos. *Trends in Neuroscience* 8, 257–266.

Bechtel, W., and A. Abrahamsen (1991). *Connectionism and the mind*. Cambridge, Mass.: Blackwell.

Benson, D. F., and N. Geschwind (1982). The alexias. In H. S. Kirshner and F. R. Freemon, eds., *The neurology of aphasia*. Lisse, Holland: Swets and Zeitlinger.

Bliss, T. V. P., and T. Lomo (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *Journal of Physiology* 232, 331–356.

Bryden, M. P. (1982). *Laterality: Functional asymmetry in the intact brain*. New York: Academic Press.

Caplan, D. (1987). *Neurolinguistics and linguistic aphasiology: An introduction*. New York: Cambridge University Press.

Caramazza, A., and M. McCloskey (1988). The case for single-patient studies. *Cognitive Neuropsychology* 5, 517–528.

Caramazza, A., and E. B. Zurif (1976). Dissociation of algorithmic and heuristic processes in language comprehension: Evidence from aphasia. *Brain and Language* 3, 572–582.

Changeux, J., and A. Danchin (1976). Selective stabilization of developing synapses as a mechanism for the specialization of neural network. *Nature* 264, 705–712.

Changeux, J., and M. Konishi (1987). *The neural and molecular bases of learning*. New York: Wiley.

Changeux, J.-P. (1985). *Neuronal man: The biology of mind* (Laurence Garey, trans.). New York: Oxford University Press.

Cherniak, C. (1990). The bounded brain: Toward quantitative neuroanatomy. *Journal of Cognitive Neuroscience* 2, 58–68.

Chugani, H. T., and M. E. Phelps (1990). Imaging human brain development with positron emission tomography. *Journal of Nuclear Medicine* 32, 23–26.

Churchland, P. S. (1986). *Neurophilosophy*. Cambridge, Mass.: MIT Press.

Churchland, P. S., and T. J. Sejnowski (1992). *The computational brain*. Cambridge, Mass.: MIT Press.

Coltheart, M., K. E. Patterson, and J. C. Marshall (1987). *Deep dyslexia*. 2nd ed. London: Routledge and Kegan Paul.

Courchesne, E., R. Yeung-Courchesne, G. A. Press, J. R. Hesselink, and T. L. Jernigan (1988). Hypoplasia of cerebellar vermal lobules VI and VII in autism. *New England Journal of Medicine* 318, 1349–1354.

Cowan, W. M. (1979). The development of the brain. In D. Flanagan, ed., *The brain*. New York: W. H. Freeman.

Cowan, W. M., J. W. Fawcett, D. D. O'Leary, and B. B. Stanfield (1984). Regressive events in neurogenesis. *Science* 225, 1258–1265.

Crick, F. H. C. (1989). The recent excitement about neural networks. *Nature* 337, 129–132.

Davidoff, J. (1991). *Cognition through color.* Cambridge, Mass.: MIT Press.

DeFries, J. C. (1990). Genetic etiology of reading disability: Evidence from a twin study. Paper presented at the meeting of the National Dyslexia Research Foundation, Rancho Mirage, Calif.

Dennis, M., and H. A. Whitaker (1976). Language acquisition following hemidecortication: Linguistic superiority of the left over the right hemisphere. *Brain and Language* 3, 404–433.

Desimone, R., and C. G. Gross (1979). Visual areas in the temporal cortex of the macaque. *Brain Research* 178, 363–380.

Diamond, M. C. (1984). Age, sex, and environmental influences. In N. Geschwind and A. M. Galaburda, eds., *Cerebral dominance.* Cambridge, Mass.: Harvard University Press.

Dudai, Y. (1989). *The neurobiology of memory.* New York: Oxford University Press.

Easter, S. S., D. Purves, P. Rakic, and N. C. Spitzer (1985). The changing view of neural specificity. *Science* 230, 507–511.

Edelman, G. M. (1987). *Neural Darwinism: The theory of neuronal group selection.* New York: Basic Books.

Ettlinger, G., C. Blakemore, A. D. Milner, and D. Wilson (1972). Agenesis of the corpus callosum. *Brain* 95, 327–346.

Feldman, R. S., and L. F. Quenzer (1984). *Fundamentals of neuropsychopharmacology.* Sunderland, Mass.: Sinauer.

Finucci, J. M., and B. Childs (1981). Are there really more dyslexic boys than girls? In A. Ansara, N. Geschwind, A. Galaburda, M. Albert, and N. Gartrell, eds., *Sex differences in dyslexia.* Towson, M. D.: Orton Dyslexia Society.

Galaburda, A. M., G. D. Rosen, and G. F. Sherman (1989). The neural origin of developmental dyslexia: Implications for medicine, neurology, and cognition. In A. M. Galaburda, ed., *From reading to neurons.* Cambridge, Mass.: MIT Press.

Galaburda, A. M., G. F. Sherman, G. D. Rosen, F. Aboitiz, and N. Geschwind (1985). Development dyslexia: Four consecutive patients with cortical anomalies. *Annals of Neurology* 18, 222–233.

Gardner, H. (1974). *The shattered mind.* New York: Vintage Books.

Gazzaniga, M. S. (1970). *The bisected brain.* New York: Appleton-Century-Crofts.

Gazzaniga, M. S. (1983). Right hemisphere language following brain bisection: A 20 year perspective. *American Psychologist* 38, 525–537.

Gazzaniga, M. S., and J. E. LeDoux (1978). *The integrated mind.* New York: Plenum Press.

Georgopoulos, A. P., and S. Grillner (1989). Visuomotor coordination in reaching and locomotion. *Science* 245, 1209–1210.

Georgopoulos, A. P., A. B. Schwartz, and R. E. Kettner (1986). Neuronal population coding of movement direction. *Science* 233, 1416–1419.

Geschwind, N., and A. M. Galaburda (1984). *Cerebral dominance: The biological foundations.* Cambridge, MA: Harvard University Press.

Geschwind, N., and N. Galaburda (1987). *Cerebral lateralization: Biological mechanisms, associations and pathology.* Cambridge, Mass.: MIT Press.

Geschwind, N., and W. Levitsky (1968). Left-right asymmetry in temporal speech region. *Science* 161, 186–187.

Gluck, M. A., and D. E. Rumelhart (1990). Neuroscience and connectionist theory. In *Developments in connectionist theory.* Hillsdale, N.J.: Lawrence Erlbaum.

Goodglass, H., and E. Kaplan (1972). *The assessment of aphasia and related disorders.* Philadelphia: Lea and Febiger.

Gordon, H. W. (1980). Right hemisphere comprehension of verbs in patients with complete forebrain commissurotomy: Use of the dichotic method and manual performance. *Brain and Language* 11, 76–86.

Graf, P., L. R. Squire, and G. Mandler (1984). The information that amnesic patients do not forget. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 10, 164–178.

Grinvald, A., E. Lieke, R. D. Frostig, C. D. Gilbert, and T. N. Wiesel (1986). Functional architecture of cortex revealed by optical imaging of intrinsic signals. *Nature* 324, 361–364.

Grossberg, S., ed. (1988). *Neural networks and natural intelligence.* Cambridge, Mass.: MIT Press.

Hanson, S. J., and C. R. Olson (1990). Connectionist modeling and brain function. In J. L. Elman, ed., *Neural network modeling and connectionism.* Cambridge, Mass.: MIT Press.

Hawkins, R. D., T. W. Abrams, T. J. Carew, and E. R. Kandel (1983). A cellular mechanism of classical conditioning in Aplysia: Activity-dependent amplification of presynaptic facilitation. *Science* 219, 400–405.

Hebb, D. O. (1949). *The organization of behavior*. New York: Wiley.

Hecht-Nielsen, R. (1990). *Neurocomputing*. Reading, Mass.: Addison-Wesley.

Hinton, G. E., J. L. McClelland, and D. E. Rumelhart (1986). Distributed representations. In Rumelhart and McClelland 1986.

Hubel, D. H. (1988). *Eye, vision, and brain*. New York: Freeman.

Hubel, D. H., and T. N. Wiesel (1977). Functional architecture of macaque visual cortex. *Proceedings of the Royal Society of London. Series* B 198, 1–59.

Hughes, J., T. W. Smith, H. W. Kosterlitz, L. A. Fothergill, B. A. Morgan, and H. R. Morris (1975). Identification of two related pentapeptides from the brain with potent opiate agonist activity. *Nature* 258, 577–581.

Hughlings-Jackson, J. (1932). Notes on the physiology and pathology of learning. In J. Taylor, ed., *Selected Writings of John Hughlings-Jackson*. London: Hodder.

Huttenlocher, P. R. (1979). Synaptic density in human frontal cortex: Developmental changes and effects of aging. *Brain Research* 163, 195–205.

Huttenlocher, P. R. (1990). Morphometric study of human cerebral cortex development. *Neuropsychologia* 28, 517–527.

Iversen, L. L. (1979). The chemistry of the brain. In D. Flanagan, ed., *The brain*. New York: W. H. Freeman.

Jenkins, W. M., and M. M. Merzenich (1987). Reorganization of neocortical representations after brain injury: A neurophysiological model of the bases of recovery from stroke. In F. J. Seil, E. Herbert, and B. M. Carlson, eds., *Progress in brain research*. New York: Elsevier.

Kandel, E. R. (1991a). Brain and behavior. In E. R. Kandel, J. H. Schwartz, and T. M. Jessell, eds., *Principles of neural science*. 3rd ed. Englewood Cliffs, N.J.: Prentice-Hall.

Kandel, E. R. (1991b). Nerve cells and behavior. In E. R. Kandel, J. H. Schwartz, and T. M. Jessell, eds., *Principles of neural science*. 3rd ed. Englewood Cliffs, N.J.: Prentice-Hall.

Kandel, E. R., and Hawkins, R. D. (1993). The biological basis of learning and individuality. In J. Piel, ed., *Mind and brain*. New York: W. H. Freeman.

Kolb, B., and Whishaw, I. Q. (1990). *Fundamentals of human neuropsychology*. New York: W. H. Freeman.

Konishi, M. (1985). Birdsong: from behavior to neuron. *Annual Review of Neuroscience* 8, 125–170.

Konishi, M. (1986). Centrally synthesized maps of sensory space. *Trends in Neuroscience* 9, 163–168.

Konishi, M. (1991). Deciphering the brain's codes. *Neural Computation* 3 (1), 1–18.

Kratz, K. E., P. D. Spear, and D. C. Smith (1976). Postcritical-period reversal of effects of monocular deprivation of striate cells in the cat. *Journal of Neurophysiology* 39, 501–511.

Lehky, S. R., and T. J. Sejnowski (1988). Network model of shape-from-shading: Neural function arises from both receptive and projective fields. *Nature* 333, 452–454.

Leuba, G., and L. J. Garey (1987). Evolution of neuronal numerical density in the developing and aging human visual cortex. *Human Neurobiology* 70, 11–18.

Lieberman, P., J. Friedman, and L. S. Feldman (1990). Syntax comprehension deficits in Parkinson's disease. *Journal of Nervous and Mental Disorders* 178, 360–365.

Linebarger, M. C., M. F. Schwartz, and E. M. Saffran (1983). Sensitivity to grammatical structure in so-called agrammatic aphasics. *Cognition* 13, 361–392.

Livingstone, M. S., and D. H. Hubel (1987). Psychophysical evidence for separate channels for the perception of form, color, movement, and depth. *Journal of Neuroscience* 7, 3416–3458.

Lynch, G., and M. Baudry (1984). The biochemistry of memory: A new specific hypothesis. *Science* 224, 1057–1063.

Macagno, E. R., U. Lopresti, and C. Levinthal (1973). Structural development of neuronal connections in isogenic organisms: Variations and similarities in the optic system of Daphnia magna. *Proceedings of the National Academy of Sciences* 70, 57–61.

Marshall, J. C. (1986). The description and interpretation of aphasic language disorder. *Neuropsychologia* 24, 5–24.

Marshall, J. C., and F. Newcombe (1981). Lexical access: A perspective from pathology. *Cognition* 10, 209–214.

McClelland, J. L., D. E. Rumelhart, and the PDP Research Group (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*. Vol. 2: *Psychological and biological models*. Cambridge, Mass.: MIT Press.

McCormick, D. A., and R. F. Thompson (1984). Cerebellum: Essential involvement in the classically conditioned eyelid response. *Science* 223, 296–299.

Merzenich, M. M. (1987). Dynamic neocortical processes and the origins of higher brain functions. In J. P. Changeux and M. Konishi, eds., *The neuronal and molecular bases of learning*. New York: Wiley.

Merzenich, M. M., P. L. Knight, and G. L. Roth (1974). Representation of cochlea within primary auditory cortex in the cat. *Journal of Endocrinology* 61, 231–249.

Milner, B. (1965). Memory disturbance after bilateral hippocampal lesions. In P. M. Milner and S. E. Glickman, eds., *Cognitive processes and the brain*. Princeton, N.J.: Van Nostrand.

Milner, B., S. Corkin, and H.-L. Teuber (1968). Further analysis of the hippocampal amnesic syndrome: 14-year follow up study of HM *Neuropsychologia* 6, 215–234.

Mishkin, M., and T. Appenzeller (1987). The anatomy of memory. *Scientific American* 256 (6), 80–89.

Nadel, L., L. A. Cooper, P. Culicover, and R. M. Harnish (1989). *Neural connections, mental computation*. Cambridge, Mass.: MIT Press.

Nauta, W. J. H., and M. Feirtag (1986). *Fundamental neuroanatomy*. New York: Freeman.

Nottebohm, F. (1970). Ontogeny of bird song. *Science* 167, 950–956.

O'Keefe, J., and L. Nadel (1978). *The hippocampus as a cognitive map*. New York: Clarendon Press.

Ojemann, G. A. (1975). The thalamus and language. *Brain and Language* 2, 1–20.

Ojemann, G. A. (1983). Brain organization for language from the perspective of electrical stimulation mapping. *Behavioral and Brain Sciences* 6, 189–230.

Oldfield, R. C. (1971). The assessment and analysis of handedness: The Edinburgh Inventory. *Neuropsychologia* 9, 97–114.

Passingham, R. E. (1982). *The human primate*. San Francisco: W. H. Freeman.

Patterson, K. E., J. C. Marshall, and M. Coltheart (1985). *Surface dyslexia*. London: Erlbaum.

Perrett, D. I., A. J. Mistlin, and A. J. Chitty (1987). Visual neurons responsive to faces. *Trends in Neuroscience* 10, 358–364.

Pinel, J. P. J. (1990). *Biopsychology*. Boston: Allyn and Bacon.

Pons, T. P., P. E. Garraghty, and M. Mishkin (1988). Lesion-induced plasticity in the second somatosensory cortex of adult macaques. *Proceedings of the National Academy of Sciences* 85, 5279–5281.

Purves, D. (1988). *A trophic theory of neuronal organization*. Cambridge, Mass.: Harvard University Press.

Rakic, P. (1979). Genetic and epigenetic determinants of local neuronal circuits in the mammalian central nervous system. In F. O. Schmitt and F. G. Worden, eds., *The neurosciences: Fourth study program*. Cambridge, Mass.: MIT Press.

Rakic, P. (1986). Mechanisms of ocular dominance segregation in the lateral geniculate nucleus: Competitive elimination hypothesis. *Trends in Neuroscience* 9, 11–15.

Rasmussen, T., and B. Milner (1977). The role of early left brain injury in determining lateralization of cerebral speech functions. *Annals of the New York Academy of Sciences* 299, 355–369.

Rolls, E. T. (1990). Spatial memory, episodic memory and neuronal network functions in the hippocampus. In L. R. Squire and E. Lindenlaub, eds., *The biology of memory*. New York: Schattauer Verlag.

Rumelhart, D. E., J. L. McClelland, and the PDP Research Group (1986). *Parallel distributed processing: Explorations in the microstructures of cognition*. Vol. 1: *Foundations*. Cambridge, Mass.: MIT Press.

Rutishauser, U., A. Acheson, A. K. Hall, D. M. Mann, and J. Sunshine (1988). The neural cell adhesion molecule (NCAM) as a regulator of cell-cell interactions. *Science* 240, 53–57.

Sacks, O. (1985). *The man who mistook his wife for a hat and other clinical tales*. New York: Summit Books.

Saul, R., and R. Sperry (1968). Absence of commissurotomy symptoms with agenesis of the corpus callosum. *Neurology* 18, 307.

Schor, R. H., A. D. Miller, and D. L. Tomko (1984). Response to head tilt in cat central vestibular neurons: I. Direction of maximum sensitivity. *Journal of Neurophysiology* 51, 136–146.

Schwartz, E. L., ed. (1990). *Computational neuroscience*. Cambridge, Mass.: MIT Press.

Scoville, W. B., and B. Milner (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery, and Psychiatry* 20, 11–21.

Sejnowski, T. J. (1986). Open questions about computation in cerebral cortex. In McClelland and Rumelhart 1986.

Sejnowski, T. J., and P. S. Churchland (1989). Brain and cognition. In M. I. Posner, ed., *Foundations of cognitive science*. Cambridge, Mass.: MIT Press.

Sejnowski, T. J., C. Koch, and P. S. Churchland (1988). Computational neuroscience. *Science* 241, 1299–1306.

Shallice, T. (1988). *From neuropsychology to mental structure*. Cambridge: Cambridge University Press.

Shepherd, G. M., ed. (1990a). The significance of real neuron architectures for neural network simulations. In Schwartz 1990.

Shepherd, G. M., ed. (1990b). *The synaptic organization of the brain.* 3rd ed. New York: Oxford.

Sherman, G. F., Rosen, G. D., and A. M. Galaburda (1989). Animal models of developmental dyslexia: Brain lateralization and cortical pathology. In A. M. Galaburda, ed., *From reading to neurons.* Cambridge, Mass.: MIT Press.

Singer, M., R. H. Nordlander, and M. Egar (1979). Axonal guidance during embryogenesis and regeneration in the spinal cord of the newt: The blueprint hypothesis of neuronal pathway patterning. *Journal of Comparative Neurology* 185, 1–22.

Skarda, C. A., and W. J. Freeman (1987). How brains make chaos in order to make sense of the world. *Behavioral and Brain Sciences* 10, 161–195.

Smith, A. (1966). Speech and other functions after left (dominant) hemispherectomy. *Journal of Neurology, Neurosurgery, and Psychiatry* 29, 467–471.

Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences* 11, 1–74.

Snyder, S. H. (1976). The dopamine hypothesis of schizophrenia: Focus on the dopamine receptor. *Journal of American Psychiatry* 133, 197–202.

Snyder, S. H. (1980). Brain peptides as neurotransmitters. *Science* 209, 976–983.

Snyder, S. H. (1986). *Drugs and the brain.* New York: W. H. Freeman.

Sperry, R. W. (1963). Chemoaffinity in the orderly growth of nerve fiber patterns and connections. *Proceedings of the National Academy of Sciences* 50, 703–710.

Sperry, R. W. (1968). Hemisphere deconnection and unity in conscious awareness. *American Psychologist* 23, 723–733.

Springer, S. P., and G. Deutsch (1989). *Left brain, right brain.* New York: W. H. Freeman.

Squire, L. R. (1987). *Memory and brain.* New York: Oxford University Press.

Steinmetz, M. A., B. C. Motter, C. J. Duffy, and V. B. Mountcastle (1987). Functional properties of parietal visual neurons: Radial organization of directionalities within the visual field. *Journal of Neuroscience* 7, 177–191.

Stephen, H., R. Bauchot, and O. J. Andy (1970). Data on the size of the brain and of various parts in insectivores and primates. In C. R. Noback and W. Montagna, eds., *The primate brain.* New York: Appleton-Century-Crofts.

Stevens, C. F. (1979). The neuron. In D. Flanagan, ed., *The brain.* San Francisco: W. H. Freeman.

Suga, N. (1990). Computation of velocity and range in the bat auditory system for echo location. In Schwartz 1990.

Swindale, N. V., M. S. Cynader, and J. Matsubara (1990). Cortical cartography: A two-dimensional view. In Schwartz 1990.

Thompson, R. (1994). *The brain.* New York: W. H. Freeman.

Trehub, A. (1991). *The cognitive brain.* Cambridge, Mass.: MIT Press.

Venezky, R. L. (1970). *The structure of English orthography.* The Hague: Mouton.

Von der Malsburg, C., and D. Willshaw (1981). Cooperativity and brain organization. *Trends in Neuroscience* 4, 80–83.

Weiskrantz, L., and E. K. Warrington (1979). Conditioning in amnesic patients. *Neuropsychologia* 17, 187–194.

Werker, J. F., J. H. V. Gilbert, K. Humphrey, and R. C. Tees (1981). Developmental aspects of cross-language speech perception. *Child Development* 2, 349–355.

Wigstrom, H., and B. Gustafsson (1985). On long-lasting potentiation in the hippocampus: A proposed mechanism for its dependence on coincident pre- and postsynaptic activity. *Acta Physiology Scandinavia* 123, 519–522.

Woolsey, C. N. (1960). Organization of cortical auditory system: A review and a synthesis. In G. L. Rasmussen and W. F. Windle, eds., *Neural mechanisms of the auditory and vestibular systems.* Springfield, Ill.: Charles C. Thomas.

Zaidel, L. (1983). A response to Gazzaniga: Language in the right hemisphere, convergent perspectives. *American Psychologist* 38, 342–346.

Zipser, D. (1990). Modeling cortical computation with backpropagation. In M. A. Gluck and D. A. Rumelhart, eds., *Neuroscience and connectionist theory.* Hillsdale, N.J.: Lawerence Erlbaum.

Zipser, D., and R. A. Andersen (1988). A back propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature* 331, 679–684.

Zipser, D., and D. A. Rumelhart (1990). The neurobiological significance of the new learning models. In Schwartz 1990.

Zola-Morgan, S., and L. R. Squire (1985). Amnesia in monkeys following lesions of the mediodorsal nucleus of the thalamus. *Annals of Neurology* 17, 558–564.

This excerpt from

Cognitive Science: An Introduction - 2nd Edition.
Neil A. Stillings, Steven E. Weisler, Christopher H. Chase,
Mark H. Feinstein, Jay L. Garfield and Edwina L. Rissland.
© 1995 The MIT Press.

is provided in screen-viewable form for personal use only by members of MIT CogNet.

Unauthorized use or dissemination of this information is expressly forbidden.

If you have any questions about this material, please contact cognetadmin@cognet.mit.edu.