from: The Computational Brain, 1992

8

Brain and Cognition

Terrence J. Sejnowski and Patricia Smith Churchland

The hope that biological and cognitive levels of investigation might be integrated has had a long history. Once it became evident that the operations of the brain were essential for thoughts and actions, discovering the biological basis for mental functions was an abiding objective (see, for example, Hippocrates (trans. 1949), Gall and Spurzheim 1968, von Helmholtz 1948, Cajal 1937, and Jackson 1958). Translating that goal into reality has been far from straightforward, however, because nervous systems are notoriously difficult to study. Until quite recently the hope often seemed frustratingly remote and unattainable.

A number of developments have improved the prospects that some of the biological mechanisms underlying cognition may be discovered in this century. In particular, important new techniques for investigating the functions of the brain have been invented in the last several decades that make possible a much more detailed structural and functional description of the brain than was previously available. Additionally we have a much better understanding of the components and dimensions of cognitive abilities (see chapter 1), and more systematic and sophisticated methods for determining the behavioral parameters of cognition have evolved (see chapter 7). Finally, the vast increases in the speed and availability of computers, together with new computational approaches to modeling in neuroscience (Sejnowski, Koch, and Churchland 1988), have allowed some problems of information processing in the nervous system to be approached effectively and have suggested new ways for thinking about cognitive operations (see chapter 4). These developments suggest that it is now possible to begin to integrate levels and to construct theories to explain functional properties of neural tissue (Churchland and Sejnowski 1988b).

The general organization of this book reflects the belief that cognitive abilities such as perception, language, memory, and attention are somehow related to brain systems. It is by no means clear, however, at what structural levels these correspondences will be made, so that the general issue of levels of organization is crucial to the enterprise from the outset. For example, it may turn out that a particular aspect of attention depends on a variety of mechanisms, some of which can be found at the level of local neural networks and others at the level of larger neural systems that reside in many different locations within the brain. Other cognitive capacities, such as planning and problem solving, may involve a complex interplay between several neural systems.

Collaborations between researchers working at different levels of investigation on common problems is often difficult. Results from one level are not always understandable or usable by researchers addressing other levels. Normally experiments are designed to address questions at a single level of organization, and the significance of the results for other levels is often not obvious. One solution is to design experiments specifically to link different levels of investigation. However, this requires an understanding of the available methodologies and knowledge of the relevant facts at both levels.

The main purpose of this chapter is to make techniques and central principles of neuroscience accessible to those who are not already immersed in the neurosciences. We do not intend to review all of the literature in neuroscience concerning the biological substrate of cognition. In particular we have not surveyed the large literature on animal behavior, which is highly pertinent. Rather, this chapter is meant to be a guide for the cognitive reader to the methods and techniques available in neuroscience for addressing cognitive issues, together with some representative examples.

8.1 Levels of Investigation

Two types of levels will be distinguished in this section—*levels of analysis* and *levels of organization.* In general it is possible to characterize a phenomenon under investigation according to its physical scale—for example, whether it is at the level of an entire perceptual system or at the level of single neurons. However, another characterization is by the level of analysis. Thus one can study a specific task performed by a particular neural structure, or one can analyze a problem in a general way that applies to all systems. In this section the relationship between these two different types of levels is explored.

Levels of Analysis

A framework for a theory of levels articulated by Marr (1982) provided an important and influential background for thinking about levels in the context of computation by nervous structures.[1] This framework drew on the concept of levels in computer science, and accordingly Marr characterized three levels: (1) the computational level of *abstract problem analysis,* decomposing the task (for example, determining structure from motion) into its main constituents; (2) the level of the algorithm, specifying a *formal procedure* to perform the task by providing the correct output for a given input; and (3) the level of *physical implementation,* constructing a working device using a particular technology.

An important element in Marr's view was that a higher level was largely independent of the levels below it, and hence computational problems of the highest level could be analyzed independently of understanding the algorithm that performs the computation. Similarly, the algorithmic problem of the second level was thought to be solvable independently of understanding its physical implementation.

Unfortunately, two very different issues were confused in the doctrine of independence. The first concerns whether, as a *matter of discovery,* one can figure out the relevant algorithm and the problem analysis independently of facts about implementation. The other concerns whether, as a *matter of formal theory,* a given algorithm that is already known to perform a task in given machine (for example, the brain) can be implemented in some other machine that has a different architecture. So far as the latter is concerned, computational theory tells us that algorithms can be run on different machines, and in that sense and that sense alone the algorithm is independent of the implementation. The formal point is straightforward: because an algorithm is formal, no specific physical parameters (for example, "vacuum tubes," "Ca^{++}") are part of the algorithm. That said, it is important to see that the purely formal point cannot speak to the issue of how best to *discover* the algorithm in fact used by a given machine, nor how best to arrive at the neurobiologically adequate task analysis. Certainly it cannot tell us that the *discovery* of the algorithms relevant to cognitive functions will be independent of a detailed understanding of the nervous system. Moreover it does not tell us that any implementation is as good as any other—and it had better not because different implementations display enormous differences in speed, size, efficiency, elegance, and so on. The formal independence of algorithm from architecture is something we can exploit to build other machines once we know how the brain works, but it is no guide to discovery if we do not know how it works.

The issue of independence of levels marks a major conceptual difference between Marr (1982) and the current generation of researchers studying neural and connectionist models. In contrast to the doctrine of independence, current research suggests that considerations of implementation are vital in the kinds of algorithms that are devised and the kind of computational insights available to the scientist. Knowledge of brain architecture, so far from being irrelevant to the project, can be the essential basis and invaluable catalyst for devising likely and powerful algorithms that have a reasonable shot at explaining how in fact the job gets done.

Levels of Organization

Marr's three-level division treats computation monolithically, as a single kind of level of analysis. Implementation and task description/decomposition are likewise each considered as a single level of analysis. Yet when we measure Marr's three *levels of analysis* against *levels of organi-*

zation in the nervous system, the fit is poor and confusing at best (Crick 1979, Churchland and Sejnowski 1988a, Shepherd 1989). To begin with, there is organized structure at different scales: molecules, synapses, neurons, networks, layers, maps, and systems. At *each* structurally specified stratum we can raise the computational question: what does that organization of elements do? What does it contribute to the wider, computational organization of the brain? In addition there are physiological levels: ion movement, channel configurations, EPSPs (excitatory postsynaptic potentials), IPSPs (inhibitory postsynaptic potentials), action potentials, evoked-response potentials, and probably other intervening levels that we have yet to learn about and that involve effects at higher anatomical levels such as networks or systems.

The range of structural organization implies, therefore, that there are many levels of implementation, and that each has its companion task description. But if there are as many types of task descriptions as there are levels of structural organization, this diversity will probably be reflected in a multiplicity of *algorithms* that characterize how the tasks are accomplished. This in turn means that the notion of *the* algorithmic level is as oversimplified as the notion of *the* implementation level.

Note also that the same level can be viewed computationally (in terms of functional role) or implementationally (in terms of the substrate for the function), depending on what questions you ask. For example, the details of how an action potential is propagated, from the point of view of communication between distant areas, might be considered an implementation because it is an all-or-none event and only its timing carries information. However, from a lower structural level—the point of view of ionic distributions—the propagating action potential is a computational construct because its regenerative and repetitive nature is a consequence of several types of nonlinear voltage-dependent ionic channels spatially distributed along an axon.

Structural Levels

Structure at every scale in the nervous system—molecules, synapses, neurons, networks, layers, maps, and systems (figure 8.1)—is separable conceptually, but not detachable physically. What is picked out as a level is actually a boundary imposed on the structure that depends on the techniques available to understand the phenomenon at hand. In the brain, they are all part of one integrated, unified biological machine. That is, the function of a neuron depends on the synapses that bring it information, and in turn the neuron processes information by virtue of its interaction with other neurons in local networks, which themselves play a particular role by virtue of their place in the overall geometry of the brain.

Accordingly, which structures really constitute a level of organization in the nervous system is an empirical, not an a priori, matter. We cannot tell in advance of studying the nervous system how many levels there

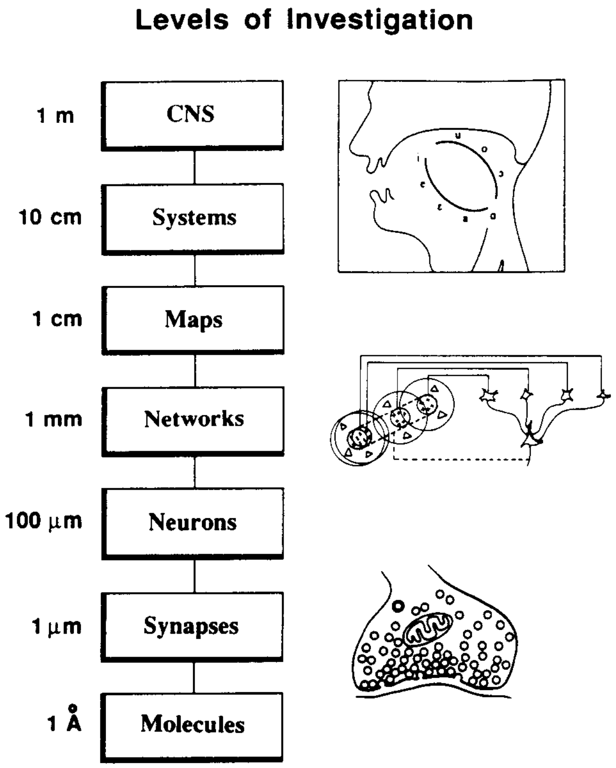

Figure 8.1 Structural levels of organization in the nervous system. The spatial scale at which anatomical organization can be identified varies over many orders of magnitude. Schematic diagrams to the right illustrate an articulatory system that produces speech sounds, a small network model for receptive fields of simple cells in visual cortex (Hubel and Wiesel 1962), and the structure of a chemical synapse (Kandel and Schwartz 1984). Relatively little is known about the properties at the network level in comparison with the detailed knowledge of synapses and the general organization of pathways in sensory and motor systems.

are nor what is the nature of the structural and functional features of any given level. The techniques that are used to study the levels are surveyed in the next section. In the next section seven general categories of structural organization are discussed (figure 8.1). The count is imprecise for several reasons. Further research may lead to the subdivision of some categories, such as "systems," into finer-grained categories, and some categories may be profoundly misdrawn and may need to be completely reconfigured. As we come to understand more about the brain and how it works, new levels of organization may be postulated. This is especially likely at higher levels where much less is known than at the lower levels.

Systems Using tract-tracing techniques, neuroanatomists have identified many *systems* in the brain. Some of the systems, such as the visual system, correspond to sensory modalities; others, such as the autonomic system, respect general functional characteristics (figures 8.2a,b). Yet others, such as the limbic system, are difficult to define and may turn

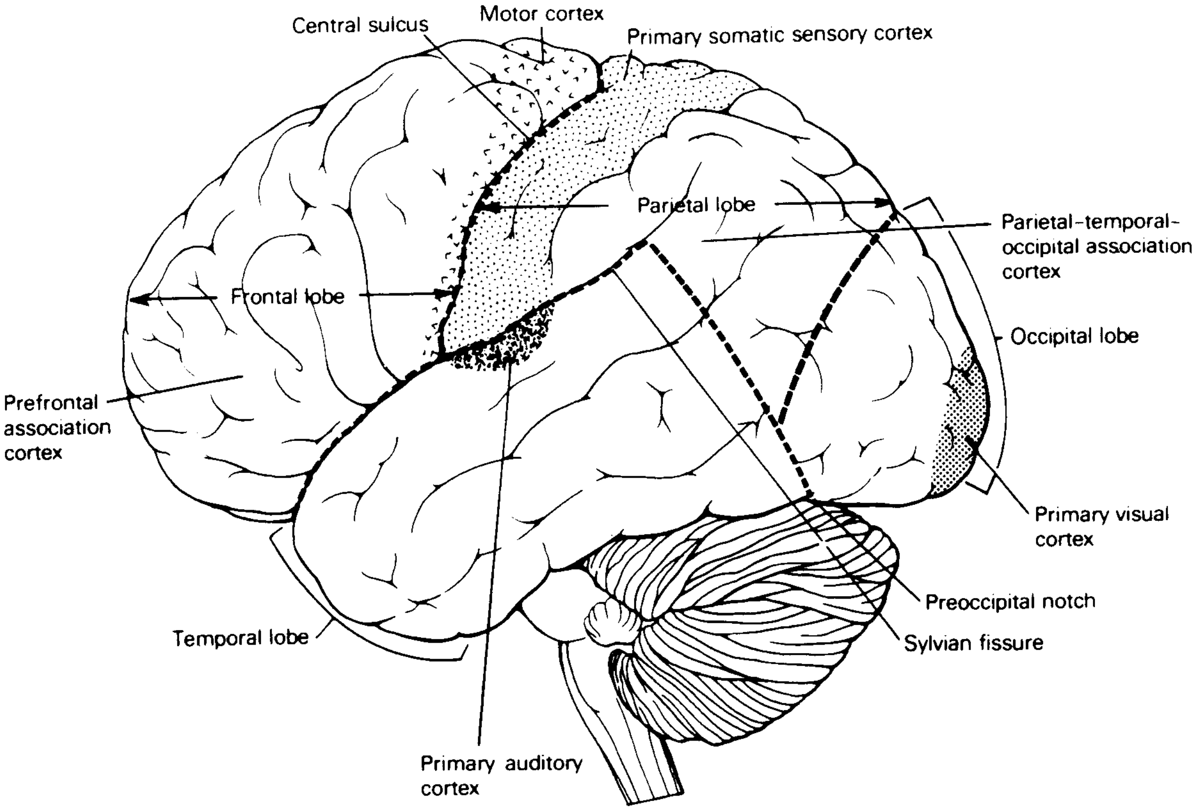

Figure 8.2a Side view of the cerebral cortex with the major divisions identified. In addition to the primary visual cortex in the occipital lobe and the primary auditory cortex near the junction of the parietal and temporal lobes, two large association areas are visible: the prefrontal-association cortex and the parietal-temporal-occipital–association cortex. The central sulcus separates the primary somatic sensory cortex from the motor cortex. The Sylvian fissure is the most prominent cleft visible in this view. (From Kandel and Schwartz 1985.)

out not to be one system with an integrated or cohesive function. The components of these systems are not neatly compartmentalized, but are distributed widely in the brain and are connected by long fiber tracts. For example, a particular brain system for long-term memory may involve such diverse structures as the hippocampus, the thalamus, the frontal cortex, and basal forebrain nuclei (Mishkin 1982). In this respect brain systems contrast quite vividly, and perhaps discouragingly, with systems designed by an engineer, in which components are discrete and functions are compartmentalized (figure 8.3).

One of the earliest systems concepts was that of a reflex arc, such as the monosynaptic reflex in the knee-jerk response (Sherrington 1906) (figure 8.4). The pathways of some reflexes have now been traced in great detail, for example, as in the vestibuloocular reflex, which stabilizes images on the retina when the head is moving (Robinson 1981), and the gill withdrawal reflex in *Aplysia,* which has been a focus for research into the molecular mechanisms of plasticity (Kandel et al. 1987). The reflex arc is not a useful prototype for brain systems in general—

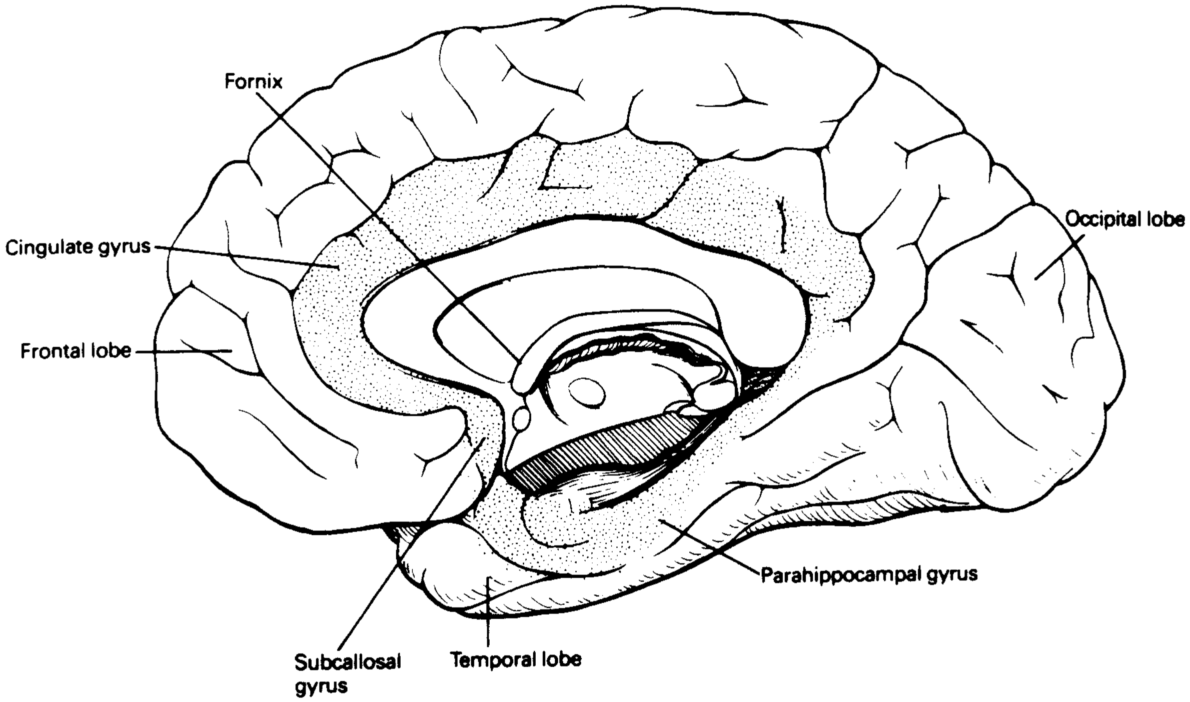

Figure 8.2b Medial view of the brain (split down the midline) showing some of the structures in the limbic system. The stippled regions of cortex that encircle the upper brainstem are primitive areas of cortex that project to the hippocampus, which is even more primitive in morphology. The hippocampus has the form of a horn and is near the arc of the fornix, which is a bundle of fibers joining the hippocampus to cortical and subcortical areas. This part of the limbic system has an important role in storing information about facts (Squire 1987). The amygdala, an olive-shaped nucleus, is located at the temporal pole of the hippocampus and has extensive connections with the hypothalamus. The emotional aspects of cortical processing are mediated by this brain region. (From Kandel and Schwartz 1985.)

or even, it appears, for most reflexes, such as the stepping reflex in the cat or the nociceptive reflex (withdrawal of limb from a painful stimulus). Take, for example, the smooth pursuit system for tracking visual targets, in which one pathway originates in the retina, leads to the lateral geniculate nucleus (LGN), to the cortex and through distinct visual topographic areas, down to the pons, and eventually to the oculomotor nuclei (Lisberger et al. 1987). Despite the machinelike quality of smooth pursuit, it is to some extent under voluntary control and depends on expectation as well as the visual stimulus. Behaviors more sophisticated than simple reflexes probably exploit more complex computational principles.

In this regard, two important features of brain systems should be mentioned. First, there are almost always reciprocal (feedback) connections between brain areas, at least as rich in number as the feedforward connections. For example, the recurrent projections from the visual cortical area V1 back to the LGN are about ten times as numerous as those from the LGN to V1. Second, although the simple models of

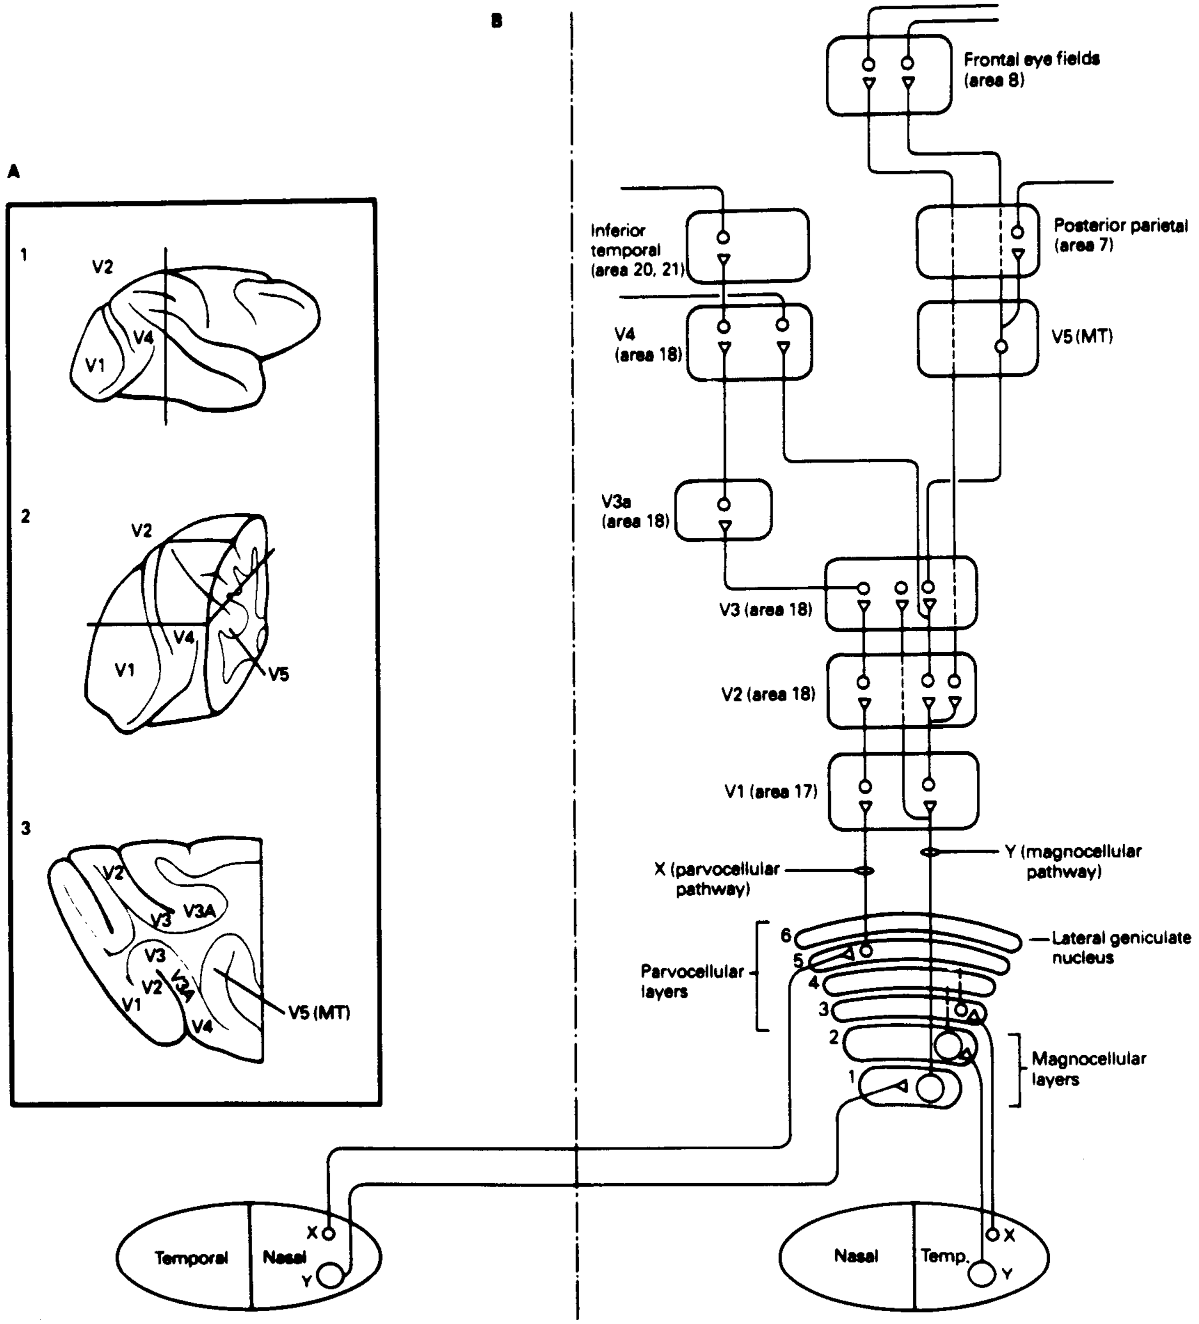

Figure 8.3 Highly schematic view of the projections from the retina to various visual areas of the cerebral cortex. (A) Visual cortical areas of the macaque monkey seen from the side (1) and through two slices, first in the coronal plane (2) and in the horizontal plane (3). The approximate location of the cortical areas V1, V2, V3, V4, and V5 are indicated (although some of these areas can be further subdivided into multiple maps of the visual field. (B) Diagram of the connections between areas in the visual pathway emphasizing three key factors in their organization. First, there are discrete levels, suggesting a hierarchy in the processing of visual information. Second, there are streams within single pathways, such as the magnocellular and parvocellular pathways. Third, there are two major divisions of higher visual processing, one following a ventral route to the inferotemporal cortex, which specializes in visual pattern recognition, and the other following a dorsal route to the posterior cortex, which specializes in processing spatial information. (From Kandel and Schwartz 1985.)

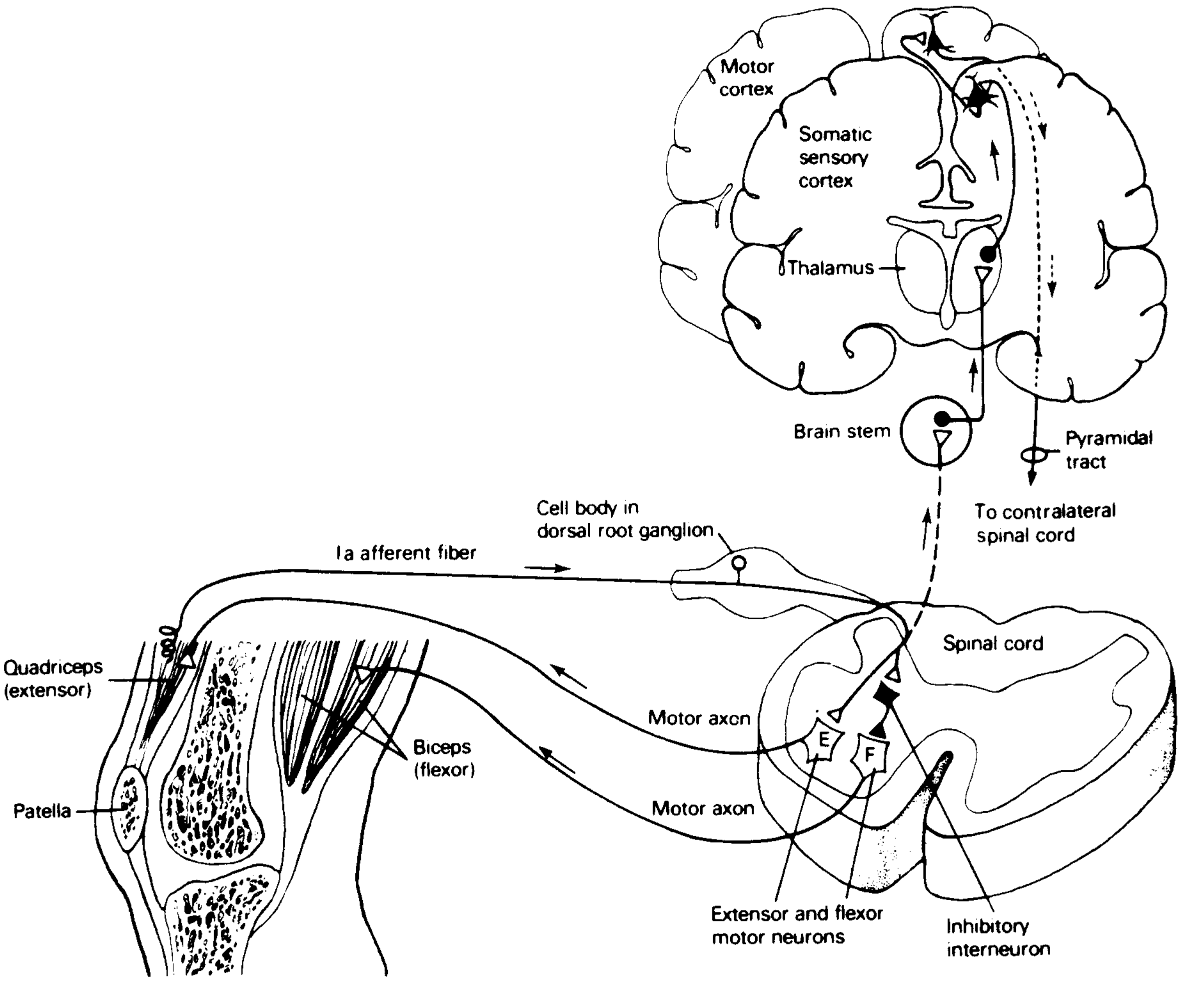

Figure 8.4 Schematic diagram of the pathways for the stretch reflex. Stretch receptors in muscle spindles react to changes in length of the muscle and afferent fibers carry this information along the dorsal roots to the spinal cord, where they synapse on extensor motoneurons, which extend the knee, and the inhibitory interneurons, which reduce activity in motoneurons that produce contractions of the antagonistic flexor muscles. Both of these actions combine to produce a coordinated expression of the knee-jerk reflex. This information is also conveyed to higher brain centers, which in turn can modify the reflex behavior through descending pathways to the spinal cord. (From Kandel and Schwartz, 1985.)

reflex arcs suggest that a single neuron may be sufficient to activate the neuron on which it synapses, in fact a large number of neurons are almost always involved, and the effect of any single neuron on the next is typically quite small. For example, an important feature in the visual system is that input from a specific neuron in the LGN generally makes relatively weak synaptic contacts on a large population of cortical cells rather than a strong synaptic effect on just one or a few neurons (Martin 1984). This implies that cortical neurons rely on a convergence of many afferents, and correlations between pairs of neurons tend to be relatively weak (Fetz and Cheney 1980, Ts'o et al. 1986). There may be interesting exceptions to this; for example, chandelier cells in cortex make inhibitory connections on the axon hillocks of their targets, and they may, as single cells, have a strong, decisive effect on their target cells. Another exception is the strong influence that single climbing fibers have on single Purkinje cells in the cerebellum.

Topographic Maps A major principle of organization within many sensory and motor systems is the topographic map. For example, neurons in visual areas of cortex, such as V1, are arranged topographically, in the sense that adjacent neurons have adjacent visual receptive fields and collectively they constitute a map of the retina. In the visual systems of monkeys physiologists have found more than twenty distinct areas, most of which are topographically mapped, though some maps are distorted and properties are unevenly represented over the surface of the map (Allman 1982, De Yoe and Van Essen 1988, Hubel and Livingstone 1987, Livingstone and Hubel 1987a,b). A similar hierarchy of multiple topographic maps are found for body location in the somatosensory system (Kaas et al. 1979; see figure 8.5), for frequency in the auditory system (Merzenich and Brugge 1973), and for muscle groups in the motor system (Ferrier 1876, Asanuma 1975). One possible exception is the olfactory system, but even odors may be spatially organized at the level of the olfactory bulb (Adrian 1953, Stewart et al. 1979). To some extent the different sensory maps can be distinguished by differences in the fine details in the laminations of neurons (see the next subsection) and their cellular properties, but often these are so subtle that only physiological techniques can distinguish boundaries between different cortical areas.

Some brainstem structures, such as the superior colliculus, also display this organization. The cerebellum appears to have patches of partial maps, though the principles do not seem clear, and these areas may not be maps in any real sense at all. Some areas seem to lack a strong topographic organization, and for other areas the topographic organization is quite complex, for example the basal ganglia (Goldman-Rakic and Selemon 1986). Cortical areas anterior to the central sulcus seem sparser in topographic maps, but research may show that what they map are abstract, not sensory, representations, and hence they cannot

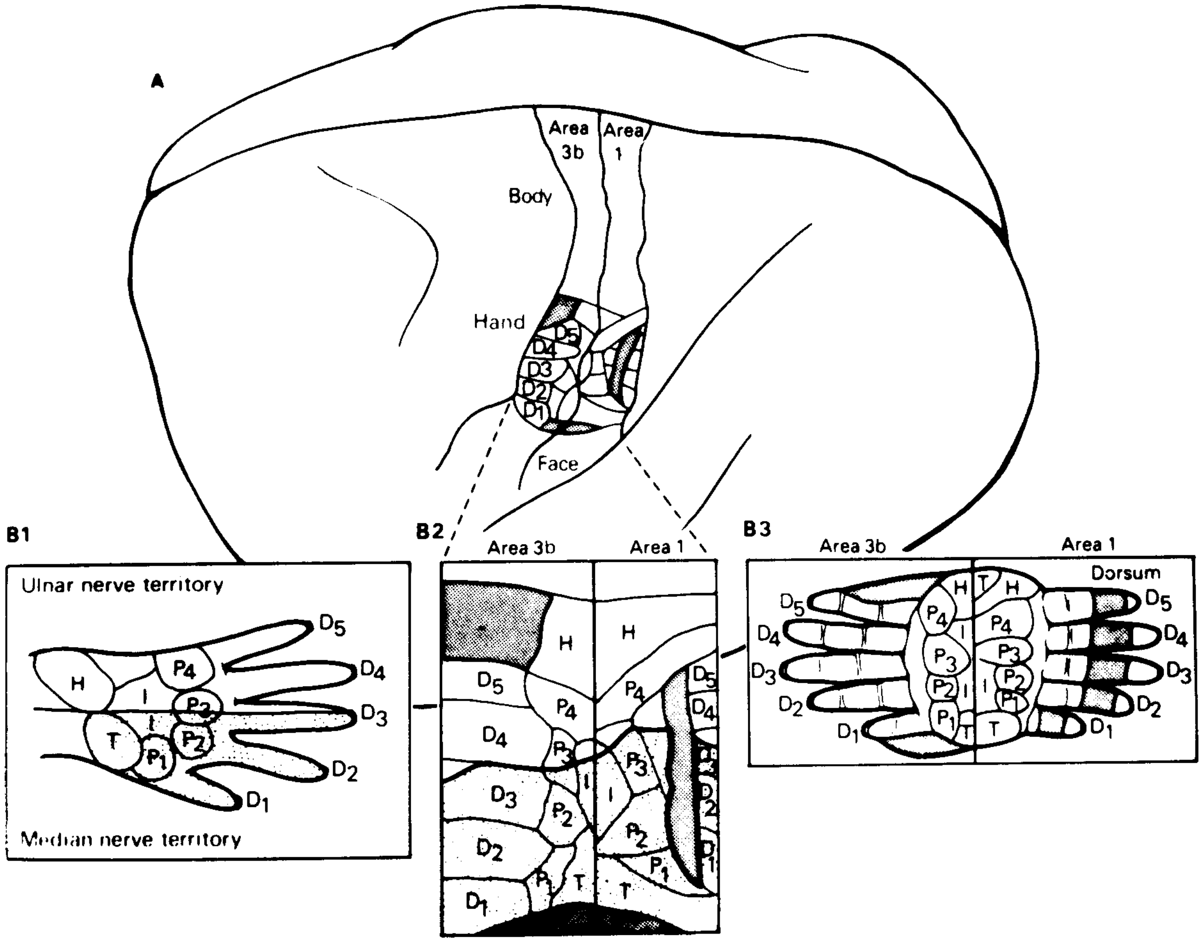

Figure 8.5 Schematic drawing of the multiple representations of the body surface in the primary somatic sensory cortex of the owl monkey. Because the cortex of the owl monkey is relative flat, most of the body representation is located on the surface rather than in the convolutions found in the species of most other primates. (A) Two representations of the hand are shown in areas 3b and 1. (B) The hand of the owl monkey is innervated by the median and ulnar nerves, which have different territory on the ventral surface (B1) and are represented in adjacent areas of cortex in each of the two maps (B2). The topographical organization of the cortical map for the ventral surface of the hand is highly ordered (B3) in both areas. Cortex devoted to the ventral surface is indicated in white; that devoted to the dorsal surface is in dark shading. D_1 to D_5 represent the five digits, P_1 to P_4 the four palmar pads, I the insular pad, H the hypothenar pads, and T the thenar pads. (From Kandel and Schwartz 1985.)

be discovered by methods used to establish response patterns to peripheral stimuli (P. M. Churchland 1985). For example, in bat auditory cortex there are topographic mappings of abstract properties such as frequency differences and time delays between emitted and received sounds, properties that may help the bat to echolocate prey (Suga 1984), and in the barn owl internal spatial maps are synthesized from binaural auditory inputs (Konishi 1986, Knudsen et al. 1987). There are some areas of cortex, such as association areas, parietal cortex, and some parts of frontal cortex, for which it has not yet been possible to find properties that form orderly mappings. Nonetheless projections between these areas remain topographic. For example, Goldman-Rakic (1987) has shown that in the monkey projections from parietal cortex to target areas in the prefrontal cortex, such as the principal sulcus, preserve the topographic order of the source neurons.

Maps of the surface of the body in the brain are formed during development by projections that become ordered, in part, through competitive interactions between adjacent fibers in the target maps (Sejnowski 1987). Some of the neurons undergo cell death during this period, and with a few rare exceptions, no new neurons are formed in the mature animal (Cowan et al. 1984). However, competitive interactions between neurons continue to some extent even in adulthood because the territory in cortex devoted to a particular part of the body surface can shift up to two millimeters, but not much farther, weeks after injury to sensory nerves or after excessive sensory stimulation (Merzenich and Kaas 1982). Thus regions in somatosensory cortex that are silenced following denervation of a sensory nerve will eventually become responsive to nearby regions of the body. It is not yet known how much of this rearrangement is due to plasticity in cerebral cortex or perhaps in subcortical structures that project to cortical maps. Nonetheless this evidence, and further evidence for synaptic plasticity, summarized below, makes it difficult to think of the machinery in the adult brain as "hardwired" or static. Rather, the brain has a remarkable ability to adapt to changes in the environment at many different structural levels and over a wide range of time scales.

Layers and Columns Many brain areas display not only topographic organization but also a laminar organization (figure 8.6a). Laminae are layers (sheets) of neurons in register with other layers, and a given lamina conforms to a highly regular pattern of where it projects to and from where it receives projections. For example, the superior colliculus receives visual input in superficial layers, and tactile and auditory input in deeper layers. Neurons in an intermediate layer of the superior colliculus represent information about eye movements. In the cerebral cortex specific sensory input from the thalamus typically projects to layer 4, the middle layer, whereas output to subcortical motor structures issues from layer 5, and intracortical projections originate mainly in

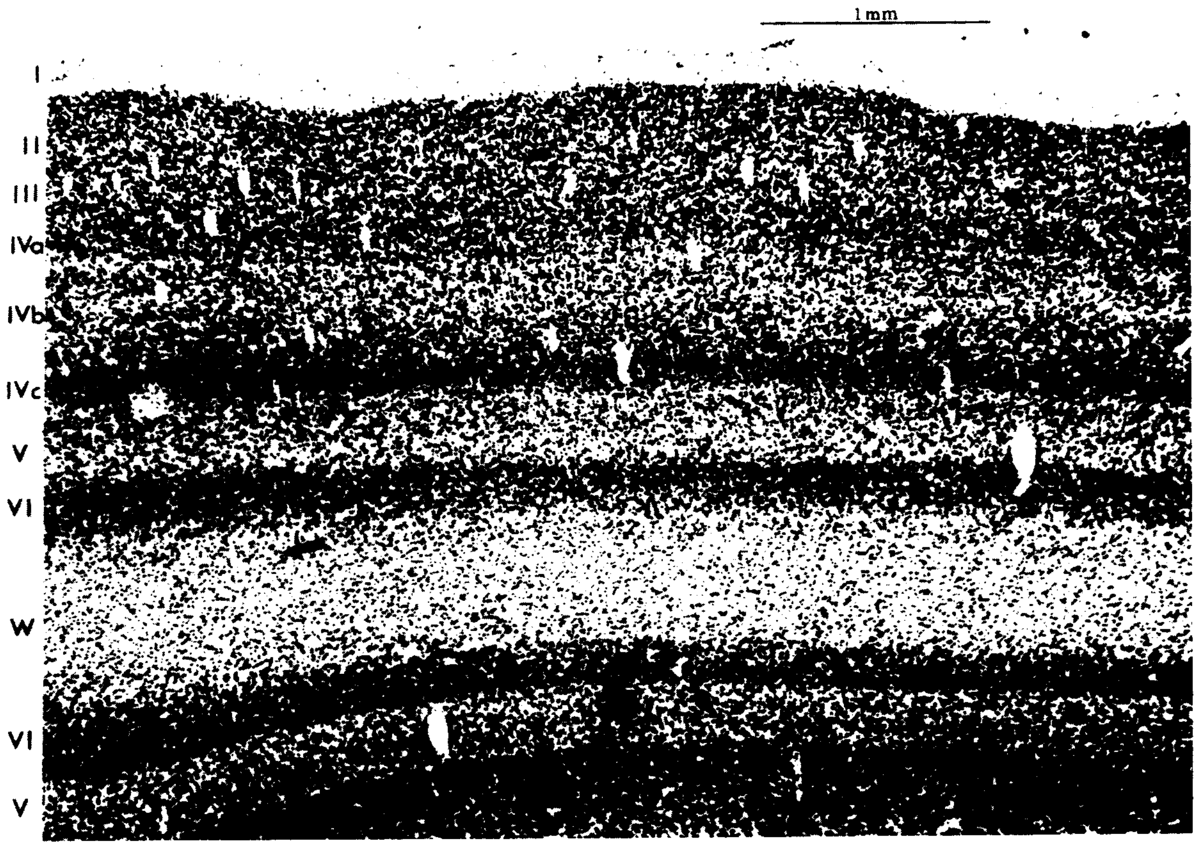

Figure 8.6a Cross-section through monkey striate cortex using cresyl violet to stain cell bodies. Laminations are clearly visible, and the layers are numbered on the left margin (W = white matter). Deeper layers of the buried fold of cortex are shown in the lower part of the figure. (From Hubel and Wiesel 1977.)

layers 2 and 3 (superficial) layers. Layer 6 mainly projects back to the thalamus. The basal ganglia do not have a laminar organization, but instead have a patchwork of islands that can be distinguished by developmental and chemical markers (Graybiel and Hickey 1982).

As well as the horizontal organization seen in laminae, cortical structures also display vertical organization. This organization consists of a high degree of commonality between cells in vertical columns, crossing laminae, and is reflected both anatomically in terms of local connections between neurons (Martin 1984, Lund 1987) and physiologically in terms of similar response properties (Hubel and Wiesel 1962). For example, a vertical penetration of an electrode in visual cortex reveals cells that share a preference for stimuli with the same orientation (for example, a bar of light oriented at about 20° from the horizontal). Another vertical penetration nearby will show cells that prefer a different orientation. Inputs and outputs are also organized columnarly, such as the eye dominance columns in V1 and inputs into the principal sulcus that alternate between parietal projections from the same side and projections from the principal sulcus in the opposite hemisphere (Goldman-Rakic 1987).

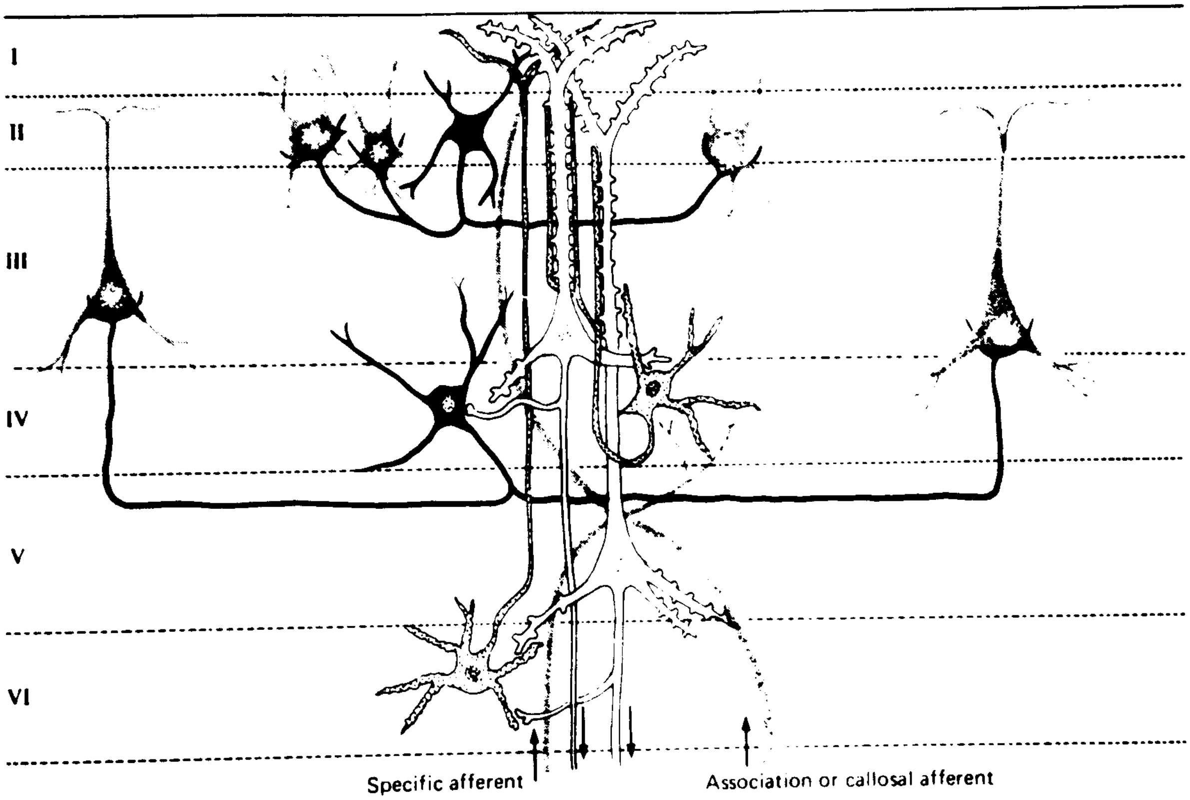

Figure 8.6b Schematic diagram of laminae and principal types of neurons in cerebral cortex. Two large pyramidal cells (white) shown in layers 3 and 5 receive synaptic contacts from the star-shaped stellate cell (stippled) in layer 4. The stellate interneuron receives specific input from principal neurons located in the thalamus. The basket cell (black) has inhibitory synaptic contacts on the somata of cortical neurons. Most of the thalamic inputs (specific afferents) are directed to cells in layer 4; association and callosal inputs from other parts of the cortex mainly terminate in the more superficial layers. (From Kandel and Schwartz 1985.)

Typically the vertically organized connectivity patterns do not result in columns with sharp boundaries, and the response properties tend to vary continuously across the cortex. Hence the expression "vertical column" may be slightly misleading. Thus for cells in visual area V1, orientation varies over the cortex smoothly, save for some fractures and singularities (Blasdel and Salama 1986), and a similar organization can be found in area V2 (Swindale et al. 1987), which receives a topographically mapped projection from V1. There are, however, places where vertical, cross-laminar columns with quite sharp boundaries are seen—for example the ocular dominance columns in layer 4 of area V1 and the "barrels" in the rodent somatosensory cortex, each of which contains cells preferentially sensitive to stimulation of a particular whisker (Woolsey and Van der Loos 1970). Sharp anatomical boundaries are, however, the exception rather than the rule. Also, the spatial scale of columnar organization can vary from about 0.3 mm for ocular dominance columns to 25 μm for orientation columns in monkey visual cortex.

Topographic mapping, columnar organization, and laminae are special cases of a more general principle, namely exploitation of geometric properties in information processing design. Spatial proximity may be an efficient way for biological systems to assemble in one place information needed to solve a problem. To consider a simple case, suppose it is necessary to compare differences between stimuli at neighboring locations, where comparison requires that signals be brought together. Then topographic organization may achieve this efficiently while minimizing the total length of the connections. This is desirable because most of the volume of the brain is already filled with axonal processes, and there are limitations on how big the brain can be as well as temporal tolerances that must be met. Lateral inhibitory interactions within the spatial maps are used to make comparisons, enhance contrast at borders, and perform automatic gain control. Mutual inhibition within a population of neurons can be used to identify the neuron with the maximum activity, a type of "winner-take-all" circuit (Feldman and Ballard 1982).

Local Networks Within a cubic millimeter of cortical tissue there are approximately 10^5 neurons and about 10^9 synapses, with the vast majority of these synapses arising from cells located within cortex (see figure 8.6b). These local networks have been very difficult to study owing to the complexity of the tangled mass of axons, synapses, and dendrites called the neuropil. Nevertheless some general features of local networks are beginning to emerge. For example, the orientation tuning of cells in V1 must emerge from nonoriented inputs and activity in local networks in ways that we are just beginning to understand (Ferster and Koch 1987).

Most of the data available on local networks are based on single-unit recordings, and to achieve a deeper understanding of the principles governing networks, it will be necessary to monitor a large population of neurons (see discussion of recording techniques in section 8.2). Even a local network involves many cells, but only small populations can be studied by exhaustive sequential recordings from single cells. Consequently we run the risk of generalizing from an atypical sample and of missing circuit properties that can be inferred only from a richer profile. Therefore, to understand the principles of local networks, much more work has to be done to determine the dynamical traffic within a larger population of cells over an extended period of time.

Computer simulations may help to interpret single-unit data by showing how a population of cells could represent properties of objects and perform coordinate transformations. For example, network models of spatial representations have been constructed that help to explain the response properties of single cells in parietal cortex (Andersen and Mountcastle 1983, Zipser and Andersen 1988; see figure 8.17). Another network model has been used to explain how the responses of single neurons in visual cortex area V4 could compute color constancy (Zeki 1983, Hurlbert and Poggio 1988). Network simulations can also suggest alternative interpretations for known response properties. For example, there are certain oriented cells in V1 whose response summates with the length of the slit or edge of light up to the borders of the receptive field, but then the response diminishes as the length increases. This property, called "end-stopping," has recently been related to the extraction of the one-dimensional curvature of contours (Dobbins et al., 1987) and the two-dimensional curvature of shapes in shaded images (Lehky and Sejnowski 1988). As we learn more about the distributed representation of information in local networks, we may achieve a better understanding of the style of computation in the brain (Sejnowski 1986, 1988, Hinton 1986).

Computers can also help in the analysis of experiments that use multielectrode recordings and optical techniques for recording. Even though the anatomical reconstruction of local networks may be prohibitive, it may be possible to reconstruct the functional properties of these networks with computer simulations (Gerstein and Aertsen 1985, Koch and Segev 1989).

Neurons Ever since Cajal's work in the late nineteenth century, the neuron has been taken as an elementary unit of processing in the nervous system (see figure 8.7). In contrast to Golgi, who believed neurons formed a continuous "reticulum" or network, Cajal argued that neurons were distinct individual cells, separated from each other by a spatial gap, and that mechanisms additional to those operating intracellularly would need to be found to explain how the signal passed from neuron to neuron. Physiological studies have borne out his judgment, though in some areas, such as the retina, syncytia of cells that are electrically coupled have been found. As it turns out, these are rather more like the structures Golgi predicted because the cells are physically joined by conducting "gap junctions." These electrical synapses are faster and more reliable than chemical transmission, but are more limited in their flexibility.

There are many different types of neurons, and different parts of the nervous system have evolved neurons with specialized properties. There are five general types of cells in the retina, for example, each with highly distinctive morphology, patterns of connectivity, physiological properties, and embryological origin. In recent years, however, physiological and chemical differences have been found within classes. For example, 23 different types of ganglion cells (whose axons project to the brain through the optic nerve) have been identified as well as 22

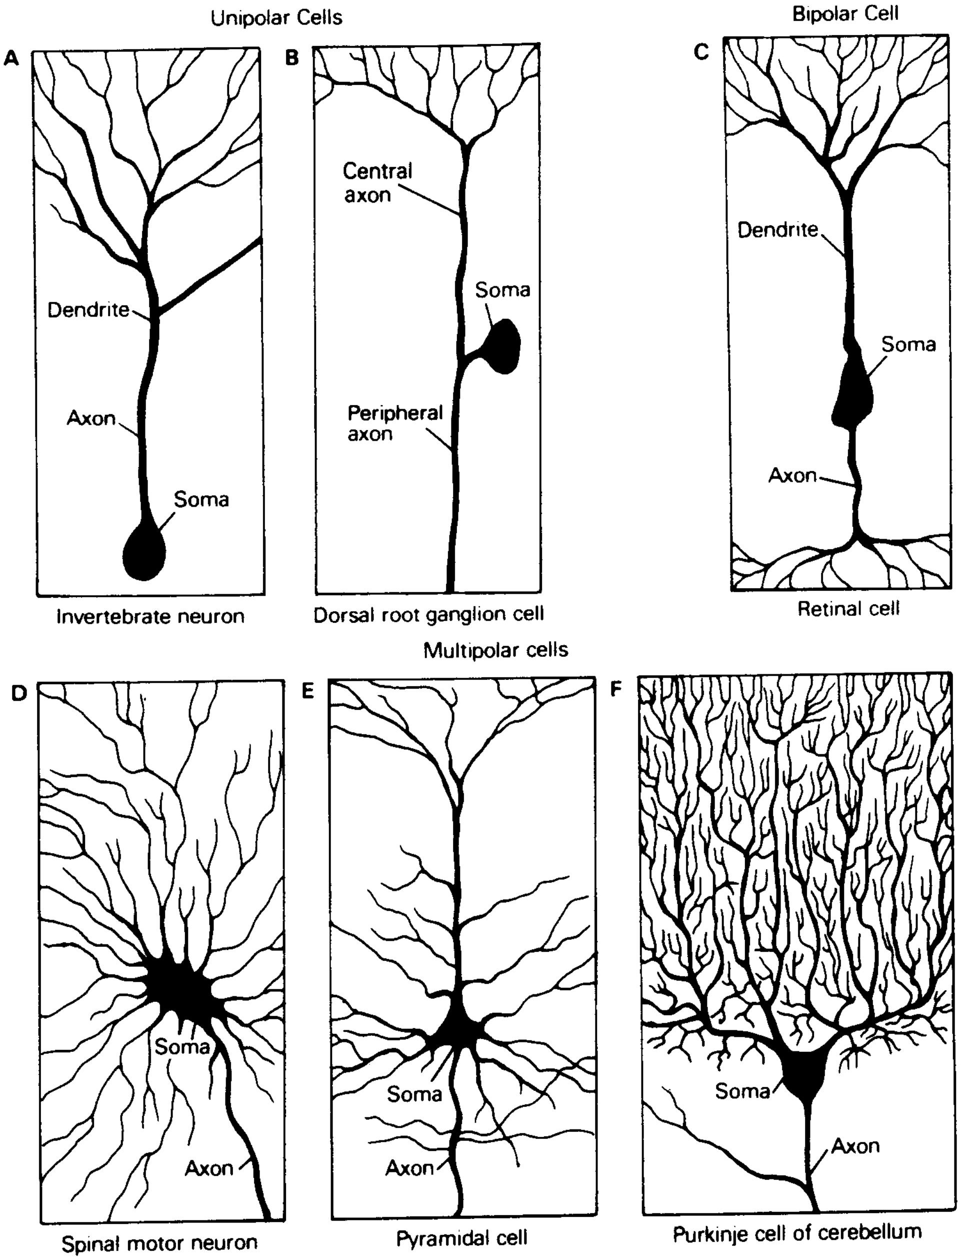

Figure 8.7 Examples of neurons illustrating the variety of shapes in different areas of the brain. The neurons are classified according to the number of processes that originate from the soma, or cell body. A unipolar cell has only a single process leaving the soma. In invertebrates (A) different segments of the unipolar process can have a receptive surface (dendrite) or transmitting regions (axons). Neurons in the dorsal root ganglion of the spinal cord (B) belong to a subclass of unipolar cells in which the soma is not directly involved with the transmission of information from the dendrites to the axon terminals. Bipolar cells (C) have two processes: the dendrite carries information toward the soma, and the axon transmits information away from the cell. Multipolar cells usually have many dendritic processes emerging from the cell body and are common in the mammalian nervous system. The spinal motoneuron (D) innervates muscle fibers and is the final common pathway for the nervous system. A pyramidal cell (E) from the hippocampus is named after the shape of the cell body. Dendrites emerge from both the apex (apical dendrites) and the base (basal dendrites). Similar pyramidal cells are found in cerebral cortex. A Purkinje cell (F) from the cerebellum has a highly branched dendritic tree that is confined to a plane. (From Kandel and Schwartz 1985.)

different types of amacrine cells (which provide lateral interactions and temporal differentiation) (Sterling 1983). There are 7 general types of neurons in the cerebellum and about 12 general types in the neocortex, with many subtypes distinguishable by their chemical properties, such as the neurotransmitters that they contain. The definition of a neuronal type is somewhat arbitrary because judgments are often made on the basis of subtle morphological differences, which can be graded rather than categorical. As more chemical markers are found, however, it is becoming clear that the diversity of neurons within cerebral cortex has been vastly underestimated. The number of subtypes of neurons in neocortex is not known, but there are at least 50 and probably less than 500 (Sereno 1988).

On the basis of their effects, neurons divide into two general classes: excitatory and inhibitory. Some neurons also have modulatory effects on other neurons, principally by releasing peptides or monoamines. Another useful classification concerns projections: some cells ramify only within an area, for example, stellate cells in cortex; and other neurons, such as pyramidal cells, have long-range projections out of an area, where the route goes via the white matter rather than directly through the cortex itself. Research on the properties of neurons shows that neurons are much more complex processing devices than previously imagined. For example, dendrites of neurons are themselves highly specialized, and some parts can probably act as independent processing units (Shepherd et al. 1985, Koch and Poggio 1987).

Synapses Chemical synapses are found in nervous systems throughout phylogeny, and they are a basic unit of structure that has been highly conserved during evolution. A synaptic bouton has a surface area of a few square microns and forms a highly stereotyped apposition with the postsynaptic membrane, which itself is highly specialized (figure 8.8). Synapses are the primary gateways by which neurons communicate with one another and consist of specialized presynaptic structures for the release of neurochemicals and postsynaptic structures for receiving and responding to those neurochemicals. Evidence is accumulating that signaling between neurons at synapses is selectively altered by experience (Alkon 1987). Nevertheless, there are other structural components of neurons that might also be modified through experience, such as the shapes and topology of dendrites as well as the spatial distribution of membrane channels (Purves and Voyvodic 1987).

The understanding of the nervous system at the subcellular level is changing rapidly, and it is becoming apparent that neurons are very dynamic and complex entities whose computational properties cannot be approximated by static response functions, a common idealization. It remains an open scientific question how the integrity of memories that span decades can remain intact if the neural substrate is as fluid

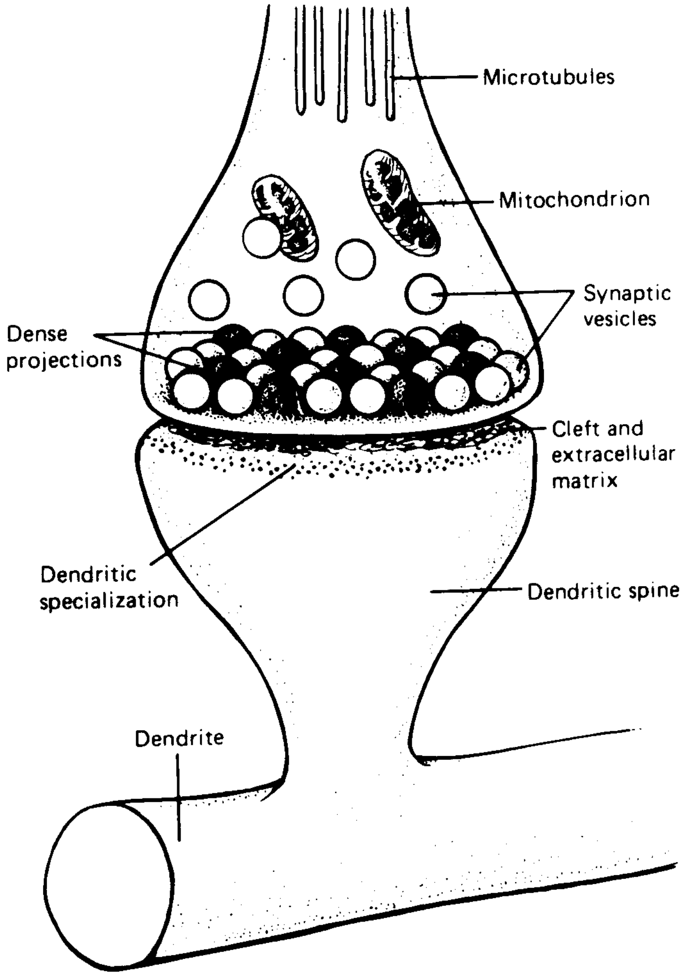

Figure 8.8 Schematic diagram of a synapse on a dendritic spine. Dense projections in the presynaptic membrane are surrounded by vesicles that presumably contain neurotransmitter molecules. This morphology characterizes a type I synapse, which is excitatory. Type II synapses (not shown) display flattened vesicles in the electron microscope after glutaraldehyde fixation and are often inhibitory. (From Kandel and Schwartz 1985.)

as preliminary reports indicate, especially if, as it seems, networks of neurons both process and store information.

Molecules The integrity of neurons and synapses depends on the properties of membranes and the internal cytoskeleton of the neuron. The membrane serves as a barrier about 10 nm thick separating the intracellular and extracellular aqueous compartments. The membrane itself is a two-dimensional fluid medium in which integral membrane proteins and other molecules form associations. Some integral membrane proteins are important in maintaining the ionic milieu inside and outside the cell. For example, membrane proteins that serve as ionic channels can be voltage-sensitive, chemically activated, or both. They may thus permit or prevent the passage of ions across the membrane, which in turn can affect the propagation of a signal down the length of the axon or neurotransmitter release at the presynaptic terminal (figure

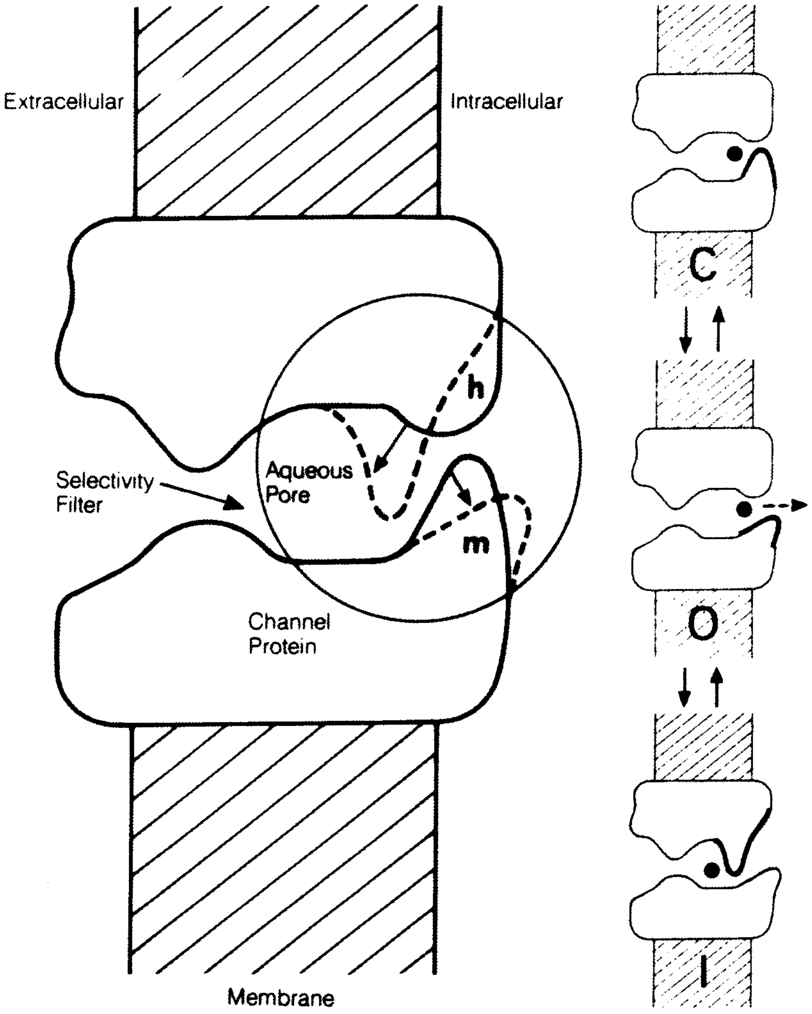

Figure 8.9 Highly schematized drawing of a sodium channel. The transmembrane protein is shown with a pore that allows sodium ions to flow between the extracellular and intracellular sides of the membrane when the gate is open. The standard Hodgkin-Huxley model of the sodium channel has two voltage-sensitive gates: The first gate (m) is normally closed at the resting potential, but opens when the cell is depolarized; the second gate (h) is normally open at the resting potential, but closes when the cell depolarizes, which inactivates the flow of ions. These two gates allow the channel to be open for only a brief period. The transitions between the closed (C) and open (O) states of the channel and the inactivated state (I) are shown on the right side of the diagram. (From Ritchie 1987.)

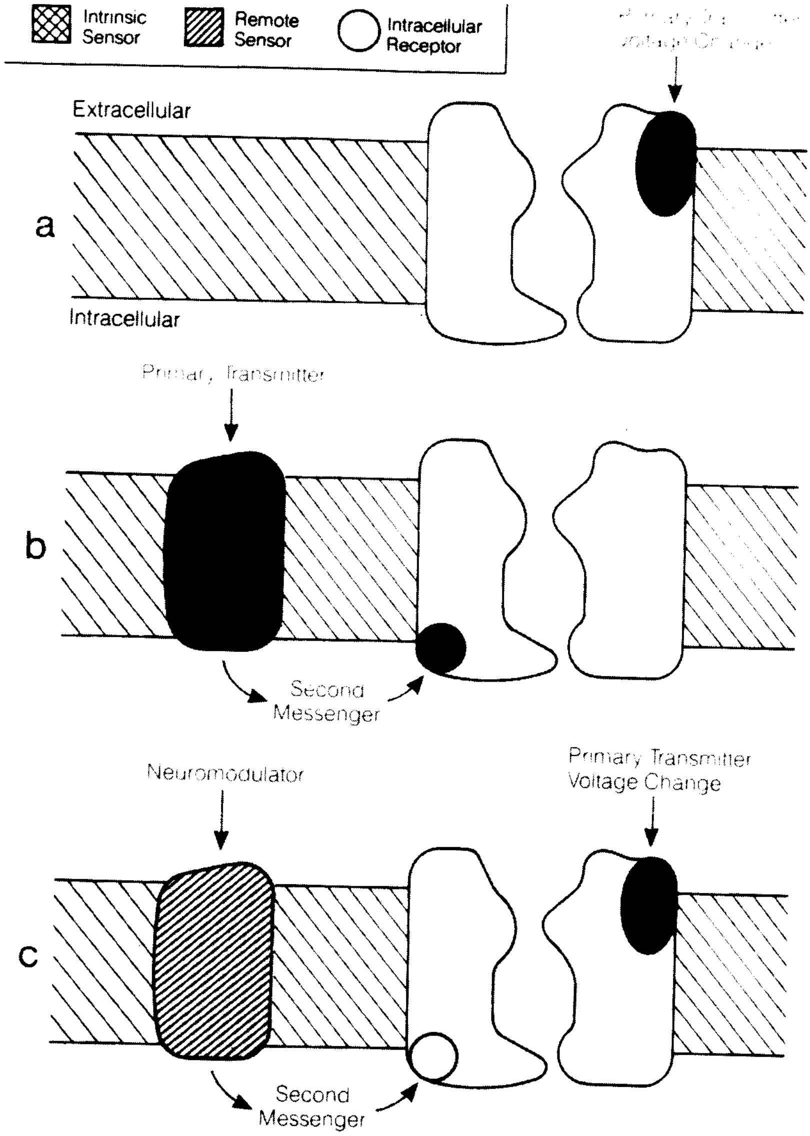

Figure 8.10 Three classes of ion channels. (a) The channel is controlled directly by voltage change or binding of a neurotransmitter molecule to a receptor. (b) The effect of the primary signal on the channel is mediated by another molecule acting as a second messenger. (c) The channel can be activated directly by the primary signal, but the response of the channel can be modulated by a neuromodular acting through a second messenger. (From Ritchie 1987.)

8.9). In a sense, the membrane allows the intracellular compartment of a neuron to respond selectively to extracellular signals, and it is this selectivity that endows different neurons with specialized information-processing capabilities.

Electrical signaling in neurons is achieved by ionic currents that are regulated by ionic channels and ionic pumps in the cell membrane. Signaling between neurons is mediated by neurotransmitter receptors in the postsynaptic membrane that respond to particular neurotransmitter molecules by transiently and selectively changing the ionic conductance of the membrane. In addition, some receptors can activate one or more second-messenger molecules that can mediate longer-term changes (figure 8.10). Second messengers in neurons can be activated by more than one receptor. Hence there is a network of interacting chemical systems within a neuron, which can itself be considered a chemical parallel-distributed processor.

8.2 Anatomical and Physiological Techniques

Nervous systems are dynamic, and physiological observations at each structural level can be arranged in a hierarchy of time scales. These scales range from microseconds, in the case of opening of single ionic channels, to days or weeks for biophysical and biochemical events underlying memory, such as long-term potentiation (McNaughton and Morris 1987, Brown et al. 1989). A large battery of techniques has been developed to try to address physiological events and processes occurring at different time scales, and by understanding what kinds of observations a given technique permits, we can begin to piece together hypotheses concerning the nature of information processing in a given structure, and how that structure contributes to the ongoing business of the brain.

Lesions

Human Studies Sometimes the brain is damaged as a result of stroke, gunshot wound, tumors, or other illnesses. Neurological assessments of the deficits and residual capacities of brain-damaged patients are an important source of information about specialization of function. When behavioral observations are correlated with a determination of the site of the lesion(s), for example, at autopsy or by magnetic resonance imaging (MRI) or positron emission tomography (PET), hypotheses can be generated concerning the brain areas that are particularly critical for certain functions. On this basis, left-hemisphere lesions have been strongly implicated in language deficits, even in most left-handed humans; bilateral lesions of the medial temporal lobe typically result in anterograde amnesia (Milner 1966, Squire et al. 1988), and lesions of the posterior parietal cortex have been implicated in loss of capacity to attend to the opposite side of the body and to the opposite hemispace (Mesulam 1985). Some lesions produce surprisingly specific perceptual and linguistic deficits (McCarthy and Warrington 1988, Damasio 1985).

Although data from clinical neurology have traditionally been and continue to be very important, they do have well-known limitations. For example, the size and exact location of lesions, which are difficult to determine, vary considerably from patient to patient. Another factor relevant in using clinical data is that patients may display some recovery of function that varies over time and that also depends on such things as the age and the gender of the patient. Additionally, the data are often difficult to interpret because a lesion may result in loss of function, not because it interferes with the specific information-processing task of a given structure, but because it interrupts fibers of passage to and from areas crucial for the function in question or some biochemical system that has a widespread influence. Moreover a lesion may affect a number of functionally distinct areas and may affect other areas through secondary degeneration. Further interpretive complications derive from other premorbid neurological and psychiatric factors, such as epilepsy, schizophrenia, and so forth. Finally, it is often difficult to find animal models for experimentation that are comparable to human cases. Despite these difficulties, important findings have been made that help to generate hypotheses concerning functional specialization of structures such as the hippocampus (see, for example, Squire, Shimamura, and Amaral 1988).

Special mention should be given in this context to split-brain studies because they involved a fairly precise surgical intervention in which any significant sparing of callosal fibers can now be detected by MRI. The disconnection effects discovered in split-brain patients (Sperry and Gazzaniga 1967) demonstrated the fragmentation of experience and awareness and confirmed the lateralization of certain functions, particularly speech production and spatioconstructive capacities, in some patients. Even in these studies, however, there are interpretive complications because all the subjects were epileptic (as opposed to normal), and there were nontrivial differences in surgical technique and in the completeness of the commissural sections. Nonetheless, studies on split-brain patients can provide a unique source of information about the global organization of processing in the brain for perceptual and cognitive phenomena such as color constancy (Land et al. 1983) and mental imagery (Kosslyn et al. 1985).

Animal Models Because humans cannot be the subjects of experimental lesions and recordings, it is essential to approach many questions concerning the human brain indirectly, via animal models. For example, the discovery in human patients that lesions to the hippocampus and related structures result in anterograde amnesia but selectively spare learning of certain skills and priming sparked the search for an animal model fitting essentially the same profile (see chapter 17). Studies on monkeys (Zola-Morgan and Squire 1984) have revealed important similarities to the human cases and, in conjunction with anatomical, pharmacological, and physiological research on the hippocampus and related structures in a variety of animals (turtles, rats, and rabbits), permit a convergence on the principles of declarative long-term memory. They also suggest new hypotheses about human memory that can be tested behaviorally. Animal studies have also been crucial to research on the neurobiological basis of sleeping and dreaming (Hobson et al. 1985).

Virtually everything we know about the microorganization of nervous systems derives from work on animal brains, and such research is absolutely indispensable if we are to have any hope of understanding the human brain. Of course there are limitations, inasmuch as there are nontrivial differences between the brains of different species, and we cannot blithely generalize from cat and monkey brains to a human brain. Even the problem of identifying homologous structures in different species can be vexing (Campbell and Hodos 1970). Nevertheless, it may be that fundamental principles can be discovered in animal models and that knowing these will provide the scaffolding for answering questions concerning those aspects of the human brain that make it unique.

Reversible Lesions and Microlesions

Some of the shortcomings of the lesion technique are being overcome by recent technical advances that make more selective intervention possible. For example, kainic acid and ibotenic acid are neurotoxic substances that destroy neurons but not fibers of passage. Additionally, the size and placement of the lesion can be carefully controlled by adjusting the amount and location of the injection. This new lesion method has been used to localize specific deficits in motion processing in MT, an extrastriate visual area in cerebral cortex (Newsome et al. 1985, Newsome and Pare 1986, Siegel and Andersen 1986).

Permanent lesions are often difficult to interpret for the reasons given above. Temporary lesions can also be made, for example, by locally cooling a region of the brain or applying local anesthetics such as lidocaine that produce measurable changes in the behavior of the animal or in the responses of neurons in other areas of the brain. Thus it is possible to separate short-term changes specific to the lesion from general or long-term alterations.

Pharmacological agents are available that can selectively interfere with particular neurons or pathways. For example, 6-hydroxydopamine, when administrated to newborn rat pups, selectively destroys all the neurons in the brain that use catecholamines, such as dopamine and norepinephrine, as neurotransmitters. Even more specific lesions are possible by taking advantage of pharmacological agents that block specific synapses. The substance 4-amino-phosphonobutyric acid (APB) has been used to selectively block a class of glutamate receptors at synapses between photoreceptors and the on-center bipolar cells in the vertebrate retina (Horton and Sherk 1984, Schiller 1982). When administered to the vitreous humor, this drug reversibly blocks the entire on-center pathway to the visual system and allows the contribution of the off-center pathway to be assessed in isolation. This is very useful because it allows us to test hypotheses about the interaction of the on-center and off-center pathways and to construct models for the origin of orientation selectivity of cells in visual cortex. Another important pharmacological agent used to investigate functional properties is aminophosphonovaleric acid (APV), which selectively blocks *N*-methyl-D-aspartic acid (NMDA) receptors. This selective blocking is fortunate because NMDA receptors are involved in the generation of long-term potentiation in the hippocampus, which is a change in the strengths of certain synapses when they are stimulated at a high rate (McNaughton and Morris 1987, Brown et al. 1989). The selective blocking of NMDA receptors may make it possible to dissect out components of some memory systems (Squire 1987).

One of the most important inhibitory neurotransmitters in the brain is gamma-aminobutyric acid (GABA). It can be applied exogenously in cerebral cortex and other areas to hyperpolarize certain neurons and thereby silence the generation of action potentials. This technique effectively lesions the cells from the local networks, and the lesion is reversible. It has been used to show that neurons in layer 6 of striate cortex contribute to the end-stop inhibition observed in the upper layers of the cortex (Boltz and Gilbert 1986). It is also possible to chronically apply GABA and look for long-term effects of the activity of neurons on the effectiveness of synapses. For example, if the eye of a kitten is sutured shut during its critical period of development, the geniculate afferents from the closed eye normally diminish in strength with the result that afterward, with both eyes open, cortical neurons respond only to stimulation of the previously open eye. Reiter and Stryker (1987) have reported that if the visual cortex of a kitten is chronically silenced by GABA and the same experimental conditions (above) are maintained, then the responses from the closed eye are preserved, whereas the responses from the open eye are abolished, reversing the normal finding. These experiments provide valuable information about the rules for synaptic plasticity (Sejnowski and Tesauro 1988).

As evidence about the composition of the brain at the molecular levels accumulates, more selective and more powerful techniques will become available for dissecting out specific neural circuits and assessing their functional significance. In particular, monoclonal antibodies, which bind specifically to particular molecules; genetic cloning techniques, which can be used to identify particular genes; and retroviruses, which can be used to insert particular genes into cells, may soon make it possible to target specific classes of cells and subclasses of synapses (Kandel 1983). Already many neurotransmitter receptors have been identified and their amino-acid sequences determined by cloning their genes. Clearly these new techniques cannot by themselves provide a deeper understanding of the function of the brain. But they can provide the answers to more detailed questions than was previously possible, though the questions must themselves evolve to exploit the potential of the techniques.

Imaging Techniques

Sherrington (1940) has described his imaginary vision of what the nervous system might look like if only the electrical activity in the brain could be seen: "Millions of flashing shuttles weave a dissolving pattern, though never an abiding one; a shifting harmony of subpatterns." With the advent of imaging techniques in the last decade, Sherrington's "enchanted loom" fantasy for large-scale visualization of the nervous system is becoming a reality, though by devices and with results that would have amazed and delighted him. Techniques for producing images of physiological activity depend on the introduction of tracers and dyes that are sensitive to physiological variables. Imaging also relies on computer power to handle the enormous flow of information, which is typically several orders of magnitude greater than that collected with traditional techniques such as single-unit recording. Some imaging techniques are noninvasive and can be safely used on a routine basis for studying normal processing in humans.

The first noninvasive mapping of brain structure was made possible by tomographic techniques that reconstructed a two-dimensional cross-section through the brain from a series of measurements on one-dimensional rays. Computed tomography (CT) uses differences in X-ray opacity of tissue as revealed on the reconstructed images to differentiate between major structures and to determine whether nervous tissue contains abnormalities. It has a spatial resolution of about 1 mm in the plane of section, which is good enough to distinguish brain regions such as the hippocampus from the amygdala. More recently MRI has been developed, and the most common maps are of hydrogen density. MRI maps have a much higher spatial resolution (about 0.1 mm in the plane of section, which is good enough to distinguish the line of Gennari, layer 4 of striate cortex), a better signal-to-noise ratio, and involve no conditions harmful to the subject, thus permitting studies of normal brains. The principle of the MRI depends on placing the tissue in strong magnetic fields and inducing and measuring changes in the magnetic orientation of the nuclei of atoms that make up the tissue. Patients are merely required to lie still in a magnetic field for about fifteen minutes. These two techniques are very useful for localizing lesions, tumors, and developmental abnormalities, but are thus far limited in not being able to assess functional damage that leaves the brain structures intact. They have provided only static images of brain anatomy, not dynamic information about brain activity. They can, however, be used in conjunction with techniques for measuring dynamic changes in brain activity (figure 8.11).

It is also possible to map other chemical elements in the brain using MRI, especially elements such as sodium and phosphorus whose concentrations vary with the functional state of the brain (Bachus et al. 1987). The concentration of sodium, phosphorus, and other chemical elements in living tissue, however, is much less than that of hydrogen, and hence the signal-to-noise ratio and the resolution with which a chemical element can be mapped using MRI are much lower.

The link between electrical activity and metabolism is exploited in the 2-deoxyglucose (2-DG) technique (Sokoloff 1984), in which a radioactively labeled sugar analog is injected into the blood and is selectively absorbed by neurons that have elevated levels of metabolic activity. In animals it is possible to section brain tissue and expose the brain tissue to X-ray film, thereby producing an image of local glucose metabolism with about a 0.1-mm resolution. Figure 8.12 is an image of responses in monkey visual cortex produced after presentation of a visual stimulus (Tootell et al. 1982), which was a flickering bull's-eye pattern. This is the first image to portray the remarkable correspondence between features in the world and activity patterns in a topographically mapped area of cortex. In humans PET can be used to image 2-DG metabolism with about a 10-mm resolution (Phelps and Mazziotta 1985).

One of the disadvantages of the 2-DG technique is that the activity must be averaged over 45 minutes, which is a very long time when we consider that a visual recognition task can be performed in under 500 ms. The time required to process the brain tissue and to expose it to X-ray film can be many months. Too much brain activity will saturate the response over large brain regions, reducing the visible differences between the areas. Moreover, it is not possible to match the experiment with a control run in the same animal because it must be killed to produce an image on film, and humans are limited to a single PET scan in one session. Because of variability between individuals, and even the same human on different sessions, averages over many subjects must be obtained. Also there is the complication of an uncertain link between glucose metabolism and electrical activity. The presumption is that they are tightly coupled in the neuropil, but although this is reasonable, it is not known with certainty in all brain areas. Nonetheless, comparisons between the 2-DG mapping and conventional electrical recordings provide a more comprehensive view of global processing (Jones, Juliano, and Whitsel 1987) and have revealed important data about topographical organization of neuronal properties within brain areas and projections between regions.

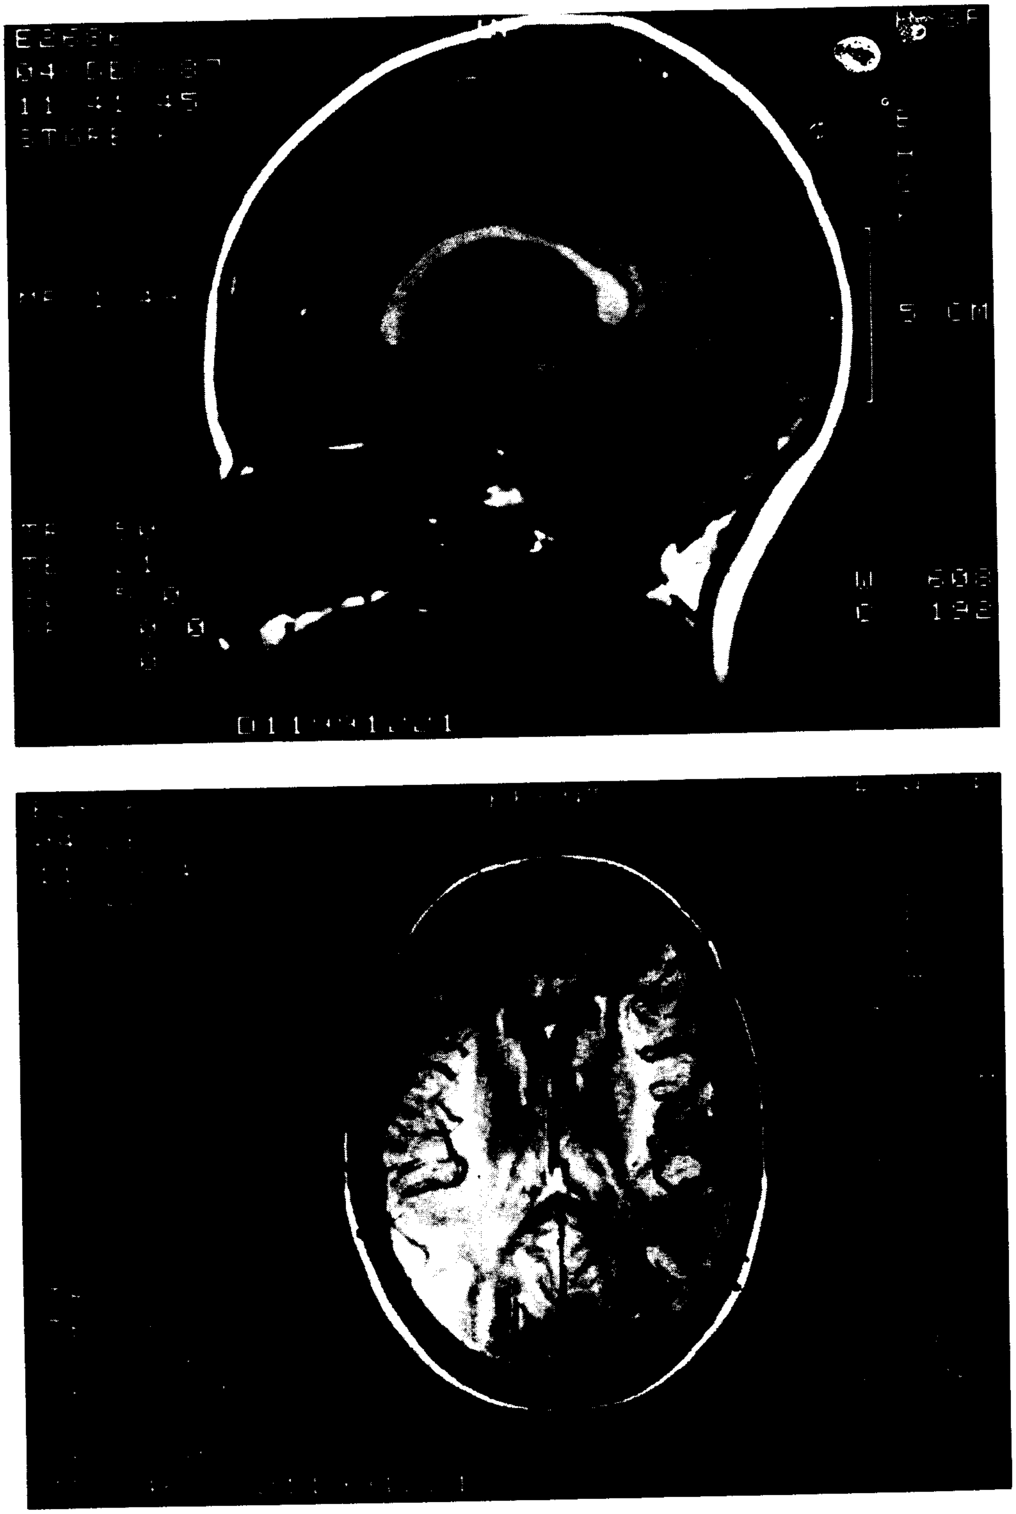

Figure 8.11 MRI (magnetic resonance imaging) of human brain. *Top*: sagittal section; *bottom*: horizontal section. (Courtesy of Bruce Crosson.)

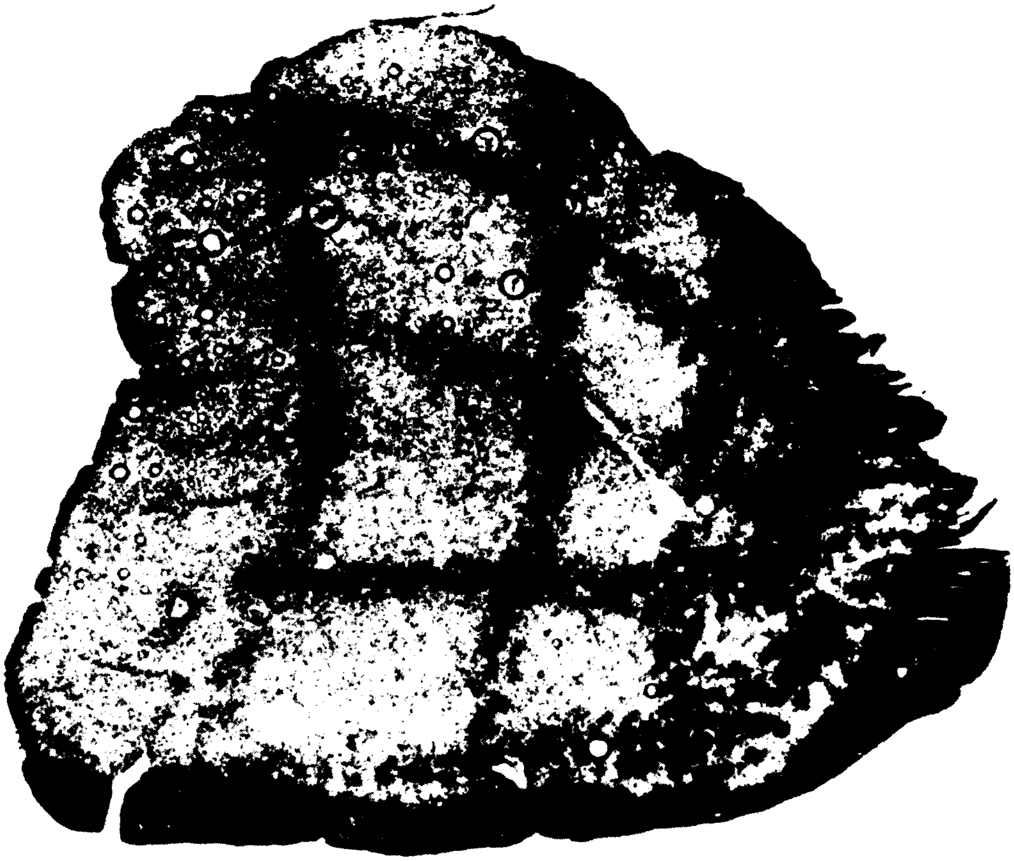

Figure 8.12 Pattern of activity in layer 4 of primary visual cortex of a monkey visualized using the 2-deoxyglucose technique. About half of the total surface area of V1 in one hemisphere can be seen in this flattened map. The visual stimulus, presented to only one eye, was a flickering target, centered on the fovea, consisting of 8 radial lines and 5 concentric circles equally spaced on a logarithmic scale. This visual cortex responds best at the onset or offset of the stimulus. The pattern of elevated brain activity is shown in this autoradiogram by the dark regions. The stripes on the dark lines are interruptions from the ocular dominance columns of the unstimulated eye. (From Tootell et al. 1982.)

Regional blood flow can be used to monitor variable metabolic demands resulting from electrical activity (Ingvar and Schwartz 1974, Roland 1984a, Raichle 1986). Blood flow is measured by following the clearance of a bolus of xenon-133 injected into the carotid artery monitored with external radiation detectors. A related method is to measure the changes in blood flow with PET following the injection of oxygen-15–labeled water into the blood. One great advantage of these methods is that several different conditions can be studied in a single session because the clearance times and half-lives are only several minutes. These techniques have been used to study voluntary motor activity (Fox et al. 1985) and selective attention to somatosensory stimuli (Roland 1984b). The current spatial resolution of PET scanning is around 10 mm, and the ultimate resolution, limited by the range of positrons, is estimated to be 2 to 3 mm. Using an averaging technique that is applicable to point sources, however, it has been possible to map the visual field in the primary visual cortex of humans with a resolution of 1 mm (Fox et al. 1987).

PET recording offers significant opportunity to investigate the localization of higher functions, including language abilities in humans. For example, cognitive tasks such as reading single words have been studied using a subtractive technique to localize individual mental operations (Petersen et al. 1988, Posner et al. 1988; see figure 8.13). There are, however, a number of confounding variables that have to be carefully separated, such as subvocal motor activity that often accompanies mental activity.

A promising approach for use on animals is the optical recording of electrical and ionic changes in the brain by direct observation. New optical dyes have been developed for noninvasive monitoring of changes in the membrane potential of neurons (Salzberg et al. 1983, Grinvald 1985, Grinvald et al. 1986). This technique has recently been used to visualize the ocular dominance and orientation columns in visual cortex (Blasdel and Salama 1986; see figure 8.14). It also appears that small changes in the absorption of red light in visual cortex can be recorded, and that these changes are correlated with electrical responses of neurons even in the absence of dyes (Grinvald et al. 1986). Ion-sensitive fluorescent dyes, such as the calcium-sensitive dye Fura-2, have also been developed that can monitor the change in intracellular ion concentration (Tsien and Poenie 1986, Connor et al. 1987). These optical techniques could be used with confocal microscopy to produce three-dimensional images of physiological activity in vivo (Boyde 1985).

Lest the beauty and remarkable achievements of these new imaging techniques inspire uncritical enthusiasm, it should be emphasized that these techniques introduce potential artifacts as well as many new problems of interpretation, and it may be some time before they can be used routinely and with confidence. Also, none of the imaging techniques is yet nearly as flexible or has as good spatial and temporal resolution *in vivo* as recording from single neurons with microelectrodes. Despite such problems, the prospects are good that we may someday obtain global views of processing in the nervous system under conditions that are close to normal.

Gross Electrical and Magnetic Recording

The earliest electrical recordings from the scalps of humans were obtained in 1929 by Hans Berger, who recorded the microvolt changes in potential with a string galvanometer. Significant differences in the electrical activity could easily be seen depending on whether the subject was awake, in deep sleep, or dreaming (Penfield and Jasper 1954; see also figure 8.15). The electroencephalogram (EEG) has also been useful in determining the general regions of the brain specialized for certain modalities and thus in locating auditory cortex, somatosensory cortex,

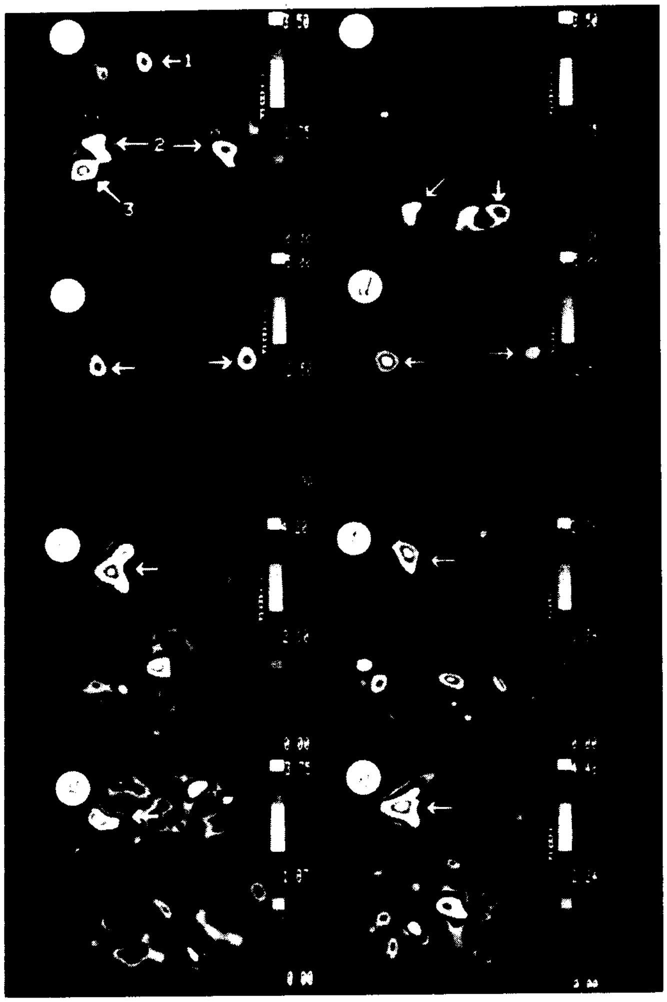

Figure 8.13 Images showing increased blood flow in regions of the brain during word processing of auditory and visual stimuli using positron emission tomography (PET). Each panel shows a horizontal slice through a human brain. The averaged difference in blood flow between two states of the brain is shown using false coloring (see scale for cerebral flow, CBF). Each of the first three pairs of panels shows tasks with auditory presentations of words at one per second (left side), and for tasks in which the words were visually presented (right side). (a, b) Average subtraction images of blood-flow change when the blood flow during rest was subtracted from blood flow during the passive response to words. The arrows indicate foci of activity in known visual and auditory sensory-processing areas. (c, d) Average subtraction images of blood-flow change when the blood flow during passive sensory presentation was subtracted from blood flow during active vocal repetition of presented words. The arrows indicate foci of activity in motor cortex. (e, f) Average subtraction images of blood-flow change when the blood flow during repetition of presented words was subtracted from blood flow during vocalization of an appropriate use for the presented word (for example, the presentation of "cake" might elicit the response "eat"). The arrows indicate strongly lateralized foci of activity in frontal cortex. (g, h) Comparison of activation on two semantic tasks. The right slice (h) is under the same condition as in e. The left slice is for change of blood flow when blood flow during passive presentation of words is subtracted from blood flow during a condition when the subject is asked to monitor the string of words for a semantic category. No motor output was required. The similar foci of activity (arrows) implicate this region of frontal cortex in semantic processing. (From Petersen et al. 1988)

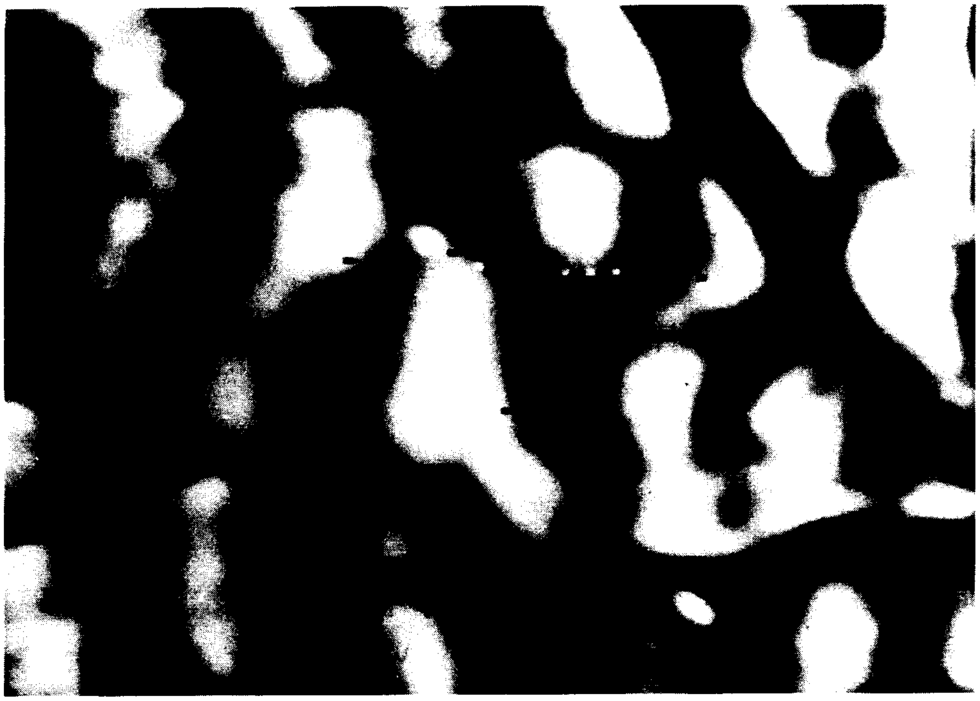

Figure 8.14 Computer-enhanced visualization of ocular dominance columns of monkey striate cortex revealed using optical recordings. A voltage-sensitive dye was applied to the cortex. One eye was closed, and the open eye was visually stimulated. Elevated activity among the neurons in alternating bands on the cortex (around 0.3 mm wide) represent ocular-dominance columns. Electrode penetrations were made into cortex tangential to layer 4, and the responses of single neurons were recorded (dots). The eye dominance recorded physiologically was consistent with the optical recording. (From Blasdel and Salama 1986.)

and so forth. A major advantage of the technique is that it is noninvasive, and it can be used on alert, behaving, normal humans.

Although the EEG has been helpful in diagnosing diseases of the brain, it has been less useful in uncovering brain mechanisms than was initially hoped. One serious difficulty in relating waveforms to underlying processing activities of neurons is that the EEG recording is a composite signal from volume conduction in many different parts of the brain, and it is far from clear what a signal means in terms of how neurons in the relevant networks are behaving. In animals, depth electrodes can be inserted and used to sort out the locations of the strong sources, but in humans this technique can be used only under special conditions where there is clear clinical justification. But even under the most favorable circumstances the localization of EEG sources is difficult and problematic.

The evoked response potential (ERP) can be extracted from the EEG recordings by signal-averaging scalp potentials that are time locked to a particular sensory stimulus or motor event, as shown in figure 8.16. (For a review see Hillyard and Picton 1987.) For example, the experimenter might present a subject with a visual stimulus, take an EEG

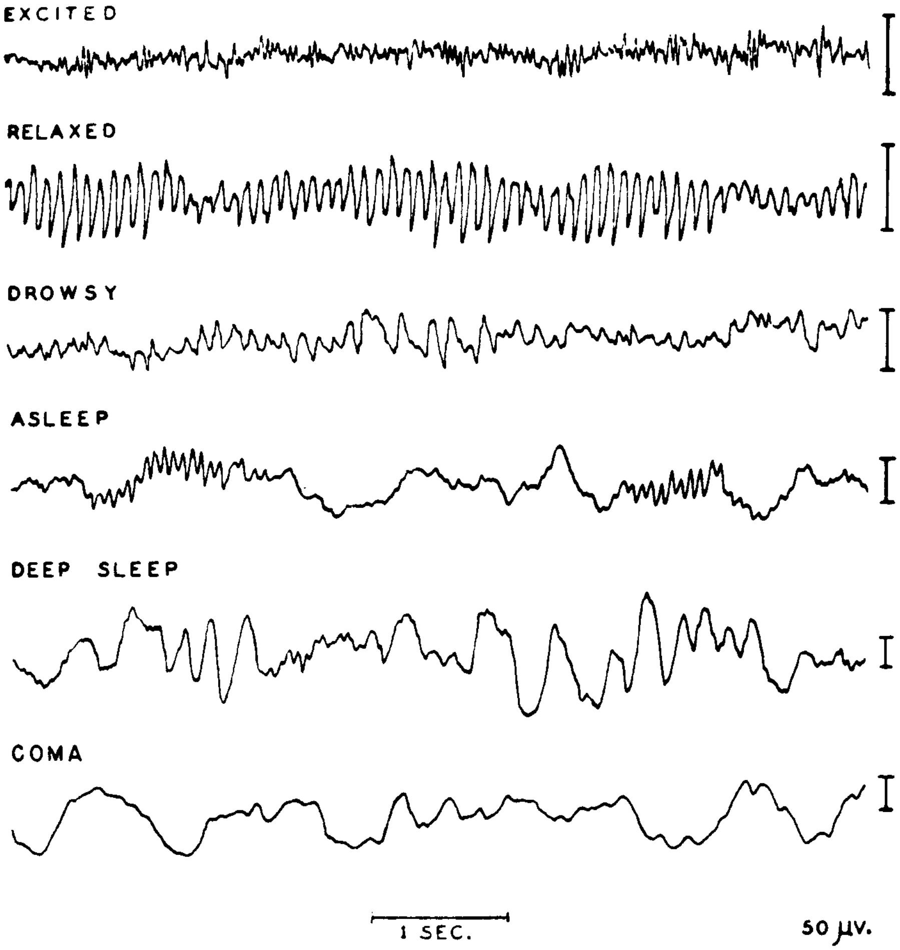

Figure 8.15 Characteristic electroencephalographic traces during variations in state of consciousness. (From Penfield and Jasper 1954.)

recording on each of ten trials, and then average those traces. Significant progress has been made in relating specific components of the ERP trace to different aspects of sensory perception. For example, reliable patterns in waveform are produced in the first 50 ms after the stimulus has been presented, and these vary reliably as a function of whether or not the stimulus was sufficiently long or intense to be conciously experienced. Other waveforms occurring later appear to reflect higher-level information processing. At about 300 ms after the stimulus, a large positive wave (P300) is produced if the stimulus was surprising or unpredicted by the subject (Donchin et al. 1983), and this has been replicated in monkeys (Paller et al. 1989). The discovery of a large negative wave (N400) under conditions where a subject is presented with a semantic incongruity (Kutas and Van Petten 1988) suggests that this technique can be used to study certain aspects of language processing. Other characteristics of the wave forms have been useful in arbitrating hypotheses about language processing and in addressing questions concerning the temporal sequence of language-processing events (Van Petten and Kutas 1987). Unfortunately, some of the com-

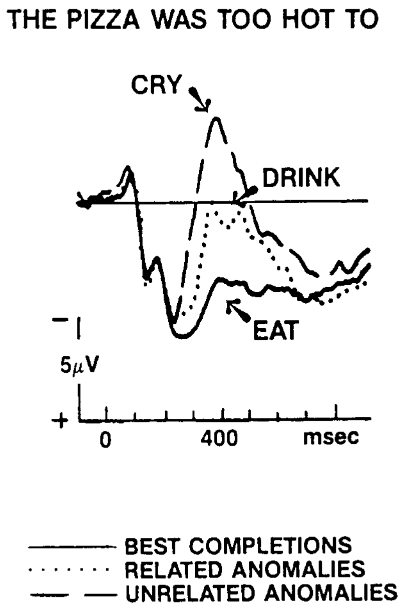

Figure 8.16 Results of an experiment using the evoked-response potential (ERP) to study language processing in humans. Ten subjects were presented with 160 different seven-word sentences, with one word presented every 700 ms via a slide projector. All subjects were presented with the same first six words (for example, *He hung it out on the line to*), but for the final word saw one of three different words (in this example, either *sing, wet,* or *dry*). Recording electrodes were placed on the scalp at the midline in the parietal zone. The figure shows the grand average ERPs to the most expected sentence completions (for example, *dry*), anomalous completions (for example, *sing*), and anomalous completions that were related to the most expected word (for example *wet*). The grand average for the final word is achieved by averaging waveforms across sentences (80 expected completions, 40 anomalous completions, and 40 anomalous but related completions) and across subjects. The N400 waveform is largest when the final word presented is an unrelated anomaly, as in *He hung it out on the line to sing,* and flattest when the final word is fitting, as in *He hung it out on the line to dry.* (Courtesy of Marta Kutas and Cyma Van Petten.)

ponents of the ERP are probably not unitary, but have several sources that have different magnitudes under different experimental conditions. For example, P300 appears to have both cortical and subcortical sources. It may be possible to model sources of evoked potentials in the brain and localize the origin of each component (Scherg and Von Cramon 1986).

The currents in neurons give rise to magnetic as well as electrical fields. These magnetic fields are not affected by volume conduction, as is the EEG, so that current sources are more easily localized than the electrical fields and can be measured with sensitive superconducting magnetometers (Williamson and Kaufman 1987). One strategy, therefore, is to try to correlate magnetic-field properties with aspects of information processing. However, it is not possible to reconstruct the internal current sources from the magnetoencephalogram (MEG) without making additional assumptions about the spatial distribution of the sources. Nonetheless, it has been possible to map parts of the visual cortex in humans and to show that primary auditory cortex is tonotopically mapped with a logarithmic frequency scale. The MEG technique is still very new and its potential has not been fully explored. In particular, arrays of magnetic sensors will soon be available that will speed up the process of producing magnetic brain maps.

Single-Unit Recording

Most of our knowledge of the response properties of single neurons has been derived from the method of single-unit recording. In this technique a microelectrode with a sharp tip is inserted into the brain and used to record local extracellular potentials. Intracellular potentials can also be recorded using extremely fine glass micropipettes; stable intracellular recordings in vivo are only possible for a few minutes, however, compared with many hours for extracellular recordings (figure 8.17). One major advantage of recording from single units is high spatial and temporal resolution, and many groundbreaking results have been achieved using the technique. For example, we have learned much about the architecture of visual cortex in anesthetized animals following the pioneering work of Hubel and Wiesel (1962). The discoveries concerning topographically mapped areas of cortex all depended on using single-unit recordings.

More recently it has become possible to study the changes in single-unit responses in the visual cortex of awake behaving animals to directed visual attention (Moran and Desimone 1985) and task-dependent variables, such as whether the animal is searching for a specific visual pattern to match a presented tactile pattern (Hanny, Maunsell, and Schiller 1988, Maunsell and Newsome 1987). The higher in the visual hierarchy one looks, however, the more difficult it is to find the adequate visual stimulus for a neuron, though there are tantalizing reports of neurons selective for hands and faces (Barlow 1985, Perrett, Mistlin, and Chitty 1987) and of neurons in the hippocampus selective for spatial location of the animal (O'Keefe and Nadel 1978). In the higher visual areas and in association cortex it is not at all clear that the stimuli chosen are the appropriate ones to use, or even how the responses should be interpreted. For example, it may someday be discovered that cells that were thought selective for faces may in fact also respond to more abstract stimuli, such as a particular class of fractal shapes (Mandelbrot 1983, Pentland 1984). Furthermore, there is always the worry that we are missing an important population of neurons because of low yield and selective sampling.

Many response properties of single neurons are highly correlated with properties of sensory stimuli and movements, but the responses of relatively few cells have been correlated with what an organism

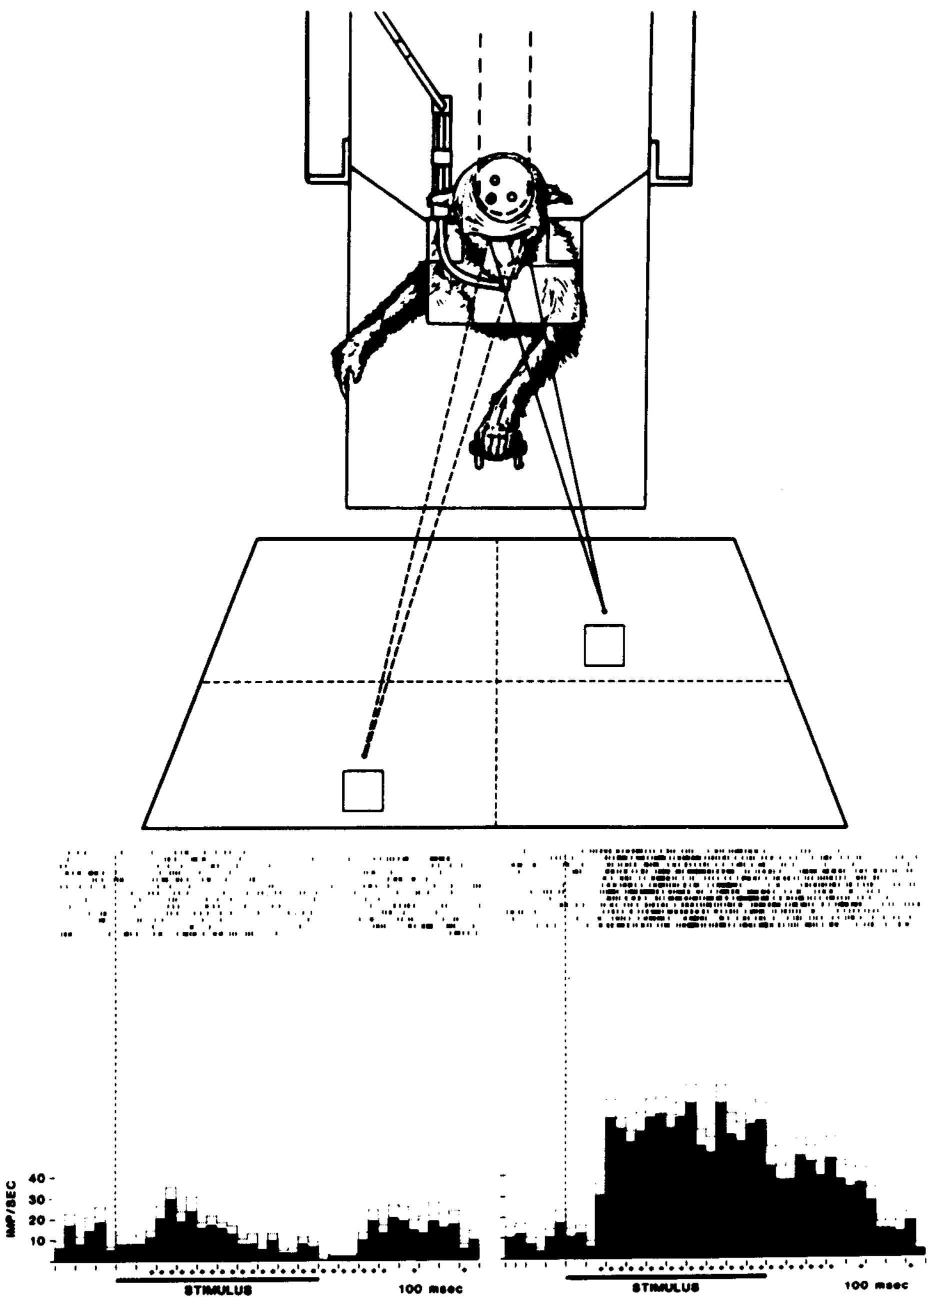

Figure 8.17 Illustration of the single-unit technique used to study the response of neurons in the parietal cortex of an awake, behaving monkey. The animal fixated a small target light placed at a series of positions on the screen, with its head fixed. The results obtained at two positions are shown here. At each fixation position a square was flashed for one second at 10° above the point of fixation. Recordings from a single neuron are shown below the screen. Each line represents a single trial and each small nick made on the line represents the discharge of an impulse by the neuron. The impulses were summed in the histograms. The right side of the figure shows the responses for the fixation to the left and down, and the responses on the left side are for fixation to the right and up. This and other experiments show that this class of neurons in parietal cortex has receptive fields that are fixed to the retina, but the degree of activation of the neuron to a visual stimulus within the receptive field is modulated by the position of the eye. Zipser and Andersen (1988) have shown that these responses can be accurately modeled by a network of neurons that transforms the location of an object from a retinal-centered coordinate system to a head-centered coordinate system. This is an example of a distributed representation of a spatial relationship. (From Andersen and Mountcastle 1983.)

perceives. For example, many neurons in our visual system respond differentially to the wavelength of light. However, the perception of color depends more on the reflectance properties of surfaces than on wavelength, so that the perception of color is roughly constant under varying illumination. Only a small subset of neurons in the visual cortex respond similarly to the perceptual report of color (Zeki 1983). If most neurons in the visual cortex respond to the spectral composition of a scene, why do we not have perceptual awareness of this information? For that matter, most neurons in V1 have an ocular preference, and some cells respond only to one eye, but when a spot of light is randomly shined into the one of the eyes, an observer cannot report which eye was stimulated despite all the information contained in single-unit responses. In general the higher a neuron in a sensory system the more likely that its response can be related to perceptual responses. For example, some neurons in V2 but not in V1 respond to illusory contours, such as Kanizsa figures (von der Heydt et al. 1984). A further problem is that responses correlated with behavior may not be causally necessary and sufficient for that behavior. For example, there is a massive change in the firing rate of hippocampal neurons in a rabbit during conditioning of an eye-blink response, but lesion of the hippocampus after training does not affect the acquired response, and rabbits without a hippocampus show normal conditioning (Berger and Thompson 1978).

Properties of networks of neurons cannot be simply inferred on the basis of the properties of small samples of cells, yet determining network properties is probably essential for understanding perceptual mechanisms. Therefore, to understand the principles of spatiotemporal coding in networks of neurons, much more work has to be done to discover what is happening in a larger population of cells. Methods for obtaining simultaneous multiunit recordings are being developed, and it is already evident that network properties to which we are blind when restricted to single-unit methods become accessible when behavior of a larger population is observed, such as synchronous firing (Llinas 1985, Reitboeck 1983, Gerstein et al. 1983). Desirable though these methods are, the technical problems in developing them are immense. Optical techniques for recording cellular responses may also prove useful in addressing population properties, but they have not yet achieved single-unit resolution in cortical structures.

A useful way to get an overview of the assorted techniques is to graph them with respect to temporal and spatial resolution. This permits us to spot spatiotemporal domains where there do not yet exist techniques to access levels of organization and to compare the strengths and weaknesses of various methods (see figure 8.18). For example, it is apparent that we lack detailed information about processing in cortical layers and columns over a wide range of time scales, from milliseconds to hours. There is also a pressing need for experimental techniques designed to address the dendritic and synaptic level of investigation in

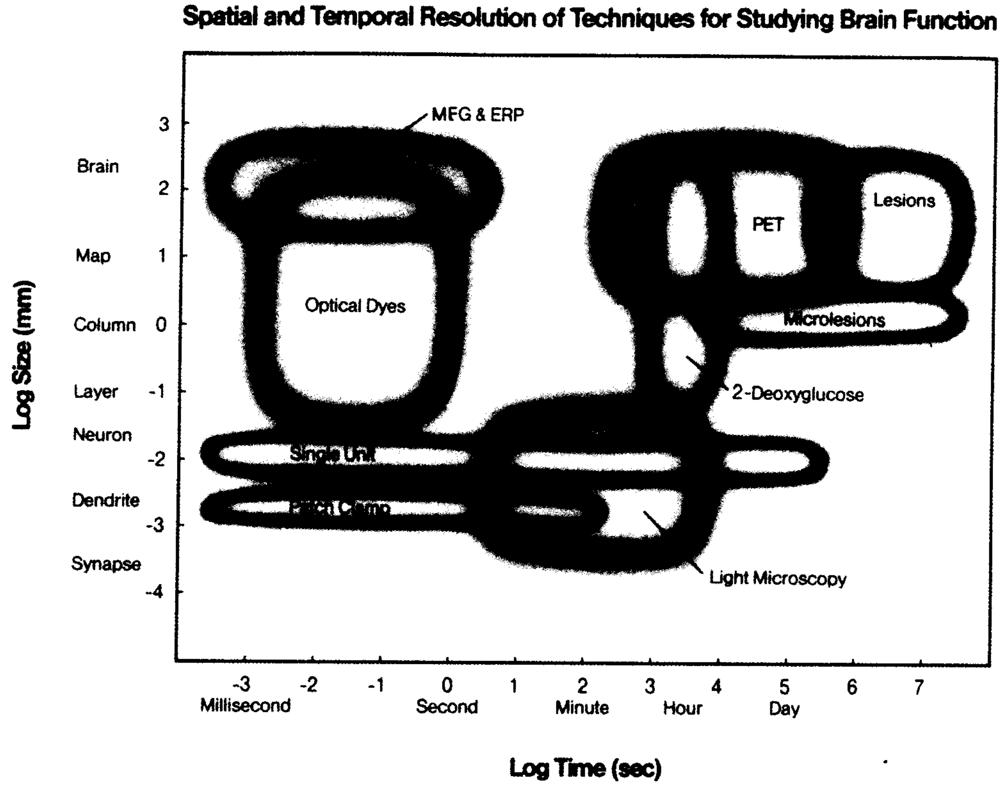

Figure 8.18 Schematic illustration of the ranges of spatial and temporal resolution of various physiological techniques for studying the function of the brain. The vertical axis represents the spatial extent of the technique, with the boundaries indicating the largest and smallest sizes of the region from which the technique can provide useful information. Thus single-unit recording can only provide information from a small region of space, typically 10 to 50 μm on a side. The horizontal axis represents the minimum and maximum time interval over which information can be collected with the technique. Thus action potentials from a single neuron can be recorded with millisecond accuracy over many hours. Patch recording allows the iconic currents through single channels to be recorded. Optical dyes have been used with cellular resolution in tissue culture, where it is possible to obtain a clear view of single cells. Recordings from the central nervous system are limited in resolution, however, by the optical properties of nervous tissue, and only about 0.1-mm resolution has been achieved. ERP (evoked-response potential) and MEG (magnetoencephalography) record the average electrical and magnetic activity over large brain regions and are limited to events that take place over about one second. The temporal resolution of PET (positron emission tomography) depends on the lifetime of the isotope being used, which ranges from minutes to an hour. It may be possible to achieve a temporal resolution of seconds with oxygen-15 to study fast changes in blood flow using temporal binning of the gamma-ray events (equivalent to the poststimulus time histogram for action potentials). The 2-deoxyglucose technique has a time resolution of about 45 minutes and a spatial resolution of 0.1 mm with large pieces of tissue and 1 μm with small pieces of tissue. Lesions allow the interruption of function to be studied both immediately after and for a long period of time following the ablation. Microlesion techniques make possible much more precise and selective interruptions of particular regions of the brain. Confocal microscopy, a promising technique for the study of nervous tissue, is a recent improvement of the light microscope for use with three-dimensional specimens. All of the boundaries here show rough regions of the spatiotemporal plane where these techniques have been used, and are not meant to indicate fundamental limitations. Purely anatomical techniques are not included. Major advances in MRI (magnetic resonance imaging) could lead to the in vivo dynamical imaging of structural modifications and chemical reaction rates (Luyten and den Hollander 1986).

cerebral cortex. Finally, every technique has its limitations, so that more than one technique may be needed at a given resolution to validate results.

8.3 General Neurobiological Constraints on Cognitive Mechanisms

A central part of the basic strategy for figuring out how a novel device works is reverse engineering. That is, when a new camera or chip appears on the market, competitors will take it apart to find out how it works. Typically, of course, competitors already know quite a lot about devices of that general kind, so the problem can be manageable. Although we have to use reverse engineering to study the brain, our starting point is much further back because we know so little about devices of that general kind. From our vantage point the brain is essentially a bit of *alien* technology, and hence it is especially difficult to know, among the available facts, which are theoretically important and which are theoretically uninteresting. We may actually misunderstand some aspects of brain organization and consequently be blocked from having some important insight into mechanisms crucial for cognition. For example, some divisions made in gross anatomy may turn out to conceal close relationships between distant brain regions. Moreover, it may turn out that the functional properties of some synapses in the central nervous system are very different from peripheral synapses in autonomic ganglia and neuromuscular junctions, on which most of our knowledge of synaptic mechanisms is based.

It would be an interesting and useful exercise to have a survey of neuroscientific opinion on this question: What are twelve basic organizational constraints on brain function from neuroscience? Or in another version: If a cognitive neuroscientist can have only twelve organizational constraints from neurobiology before being locked in a room to develop functional hypotheses, which twelve would they be? Short of having formally conducted such a survey, we conjecture that the following dozen are among those likely to be on such a list, although we recognize that opinion can diverge in considerable and surprising ways, and also that any current list would undoubtedly be quickly outdated (for comparable lists see Shepherd 1988 and Crick and Asanuma 1986):

1. It is estimated that in the human nervous system there are about 10^{12} neurons and about 10^{15} synapses.

2. An action potential lasts about 1 ms; synaptic transmission, including electronic conduction in dendrites, takes about 5 ms. Transmission velocity in myelinated axons is about 10 to 100 m/s; in unmyelinated axons it is less than 1 m/s. These are general ranges, not precise values.

3. Synapses are excitatory or inhibitory and synaptic potentials can last from a millisecond to many minutes (Kuffler 1980). The integration of synaptic potentials can be additive or multiplicative depending on the ionic conductances and the synaptic geometry. Certain neurotrans-

mitters, such as some peptides and monoamines, act at several tenths of a millimeter from their release sites, though with a delay and a long-lasting effect (Kuffler 1980). It has also been discovered that at some synapses the same transmitter can activate several different receptors on the postsynaptic membrane and different neurotransmitters can be released from a single bouton under different stimulus conditions (Kuffler 1980). Peptide neurotransmitters, for example, almost always co-localize with at least one other neurotransmitter (Hokfelt 1987). Hormones in the general circulation can alter neural activity.

4. Typically, cell-to-cell interactions are weak (amounting to 1 to 5 percent of the firing threshold), and cells receive inputs from thousands of other cells. There may be some interesting exceptions to this, such as the chandelier cells in cerebral cortex (Martin 1984).

5. The brain appears to be highly parallel in that there are many parallel streams of input for a given function. For example, in the monkey two parallel streams from the retina, starting with different types of ganglion cells, project to two distinct sets of layers of the lateral geniculate nucleus, the parvocellular and magnocellular layers respectively, which in turn project to distinct sublaminae in layer 4 of cortical area V1 of the visual cortex (Hubel and Livingstone 1987; Livingstone and Hubel 1987a,b). The system also appears to be hierarchical, insofar as there are multiple feedforward stages. The latency of response to a sensory stimulus, such as a flashing light, reveals a temporal hierarchy of processing within each processing stream.

6. Not everything is directly connected to everything else, and connectivity appears to be highly specific. Most connections are between, not within, cell classes (Sereno 1988). Within certain limits there is considerable plasticity in connectivity, as demonstrated by the changes in ocularity of cortical cells of kittens induced by keeping one of the eyes closed during the critical period of development.

7. Cortical layers in sensory areas display considerable regularity in projection profiles: feedforward connections are mainly from layer 4, feedback connections are mainly non-4, namely 2, 3, and 5, and are made via white matter (Maunsell and Van Essen 1983). Pyramidal cells in layers 5 and 6 mainly project to subcortical structures such as the basal ganglia, superior colliculus, claustrum and thalamus (see figure 8.6b). Cortical connections from one topographically mapped area to another are strictly excitatory.

8. In addition to the specific system projecting to the neocortex via the thalamus, such as is seen in the visual, auditory, and somatosensory systems, there are five sources of widely projecting neurons, each associated with a specific neurotransmitter, which may be important in the sleep-dream-wake cycle, in memory, and in awareness and attention. The five sources are the locus coeruleus in the brain stem (norepinephrine), the raphe nucleus in the midbrain (serotonin), the substantia

nigra in the midbrain (dopamine), the nucleus basalis in the basal forebrain (acetylcholine), and special groups of cells in the mammillary region of the hypothalamus (GABA). The neurotransmitter receptors for these systems belong to a superfamily of receptors whose effects are mediated by second-messenger effects (see figure 8.10). The effects are longer lasting than those produced by neurotransmitters acting on receptors from other superfamilies, such as the class of glutamate receptors.

9. Under the classical definition, the receptive field of a cell is that region of the sensory field from which an adequate sensory stimulus will elicit a response. Receptive field properties in higher areas of visual cortex differ from the early stages, where receptive fields are quite small (one-sixth of a degree in the foveal region of V1), and cells are tuned to respond to quite simple stimuli (spot of light, bar of light). In inferotemporal cortex, for example, the fields are much larger (10° to the whole visual field), and the stimuli evoking responses are more complex (a hand, a face).

10. Neurons are rather coarsely tuned, and the receptive fields of cells overlap. This suggests that precision is achieved through overlapping redundancy and population coding. It is not known to what extent particular objects and single items are coded by small sets of neurons ("grandmother cells"). It is not generally possible to characterize correctly the tasks of neurons in a particular area, even when something is known about the response properties from single cells in that area. More precisely, models show that the *function* of a sensory neuron depends on its output "projective" field as well as its receptive field (Lehky and Sejnowski 1988).

11. Events outside the classical receptive field of a cell have been found to selectively modulate the responses of the cell (Allman et al. 1985). The effects are selective since they vary as a function of the type of surrounding stimuli. For example, certain wavelength-dependent neurons in V4 are influenced by the color balance in the surroundings (Zeki 1983). The surrounding effects of cells in the middle temporal area (MT), where receptive fields are typically 5° to 10°, can extend from 40° to 80°.

12. Last, but of fundamental importance, the nervous system is the product of a long evolution. The original functions of some parts have been altered by layer upon layer of modifications. For example, our ability to mentally image complex objects depends on visual processing systems that probably evolved for other functions (Kosslyn 1987, 1988). We cannot expect the design of the brain to resemble anything that a human would consider optimal. Much of our intuition must be guided therefore by biological facts rather than by logical assumptions. This may doom forever someone who wants to understand brain function but does not have access to more of the relevant experimental literature.

8.4 Models

While we need experimental data concerning the properties of neurons, we also need to find models that explain how neural networks manage to represent such things as surfaces, optical flow, and color constancy; how networks learn, store, and retrieve information; how they accomplish sensorimotor integration; and so forth. Ideally, modeling and experimental research have a symbiotic relationship, such that each informs, corrects, and inspires the other.

To treat a mechanism as a black box is equivalent to deciding to ignore certain details of the mechanism in order to find the properties that play a prominent role in the function. Such simplifying and idealizing is essential to theorizing in science in general, though it is notoriously difficult to determine which properties one can safely ignore in constructing the theory of the function, and certainly no decision procedure exists for that problem.

We could treat the brain as a whole as a black box, but the trouble with this strategy is that there are indefinitely many computational models one might dream up, and none might be even close to how the brain in fact achieves solutions to difficult computational problems. The more profitable strategy would be to treat the brain as a lot of little black boxes. For example, one could take an individual neuron as a black box, thereby deciding not to worry for the nonce about such matters as types of single channels in neuronal membranes, the details of dendritic responsivity, the role of microtubules, and so on. One could choose a slightly larger size of black box, taking, say, cortical columns as units, where the details of connectivity within the column are ignored in the model but intercolumn connectivity is analyzed and incorporated. Or, going larger yet again, one could treat the visual areas as wholes, or the hippocampus or cerebellar cortex as the level of organization one tries to address.

Although diverse kinds of things are presented as models for some part of the nervous system, it is useful to distinguish between ideas that are genuinely and strongly predictive of some aspect of nervous system dynamics or anatomy and ideas that, though not so predictive, demonstrate that the nervous system *could be* governed by the principles specified in the idea. The first—genuinely predictive hypotheses—we call *models*, and the second we call *demonstrations.*[2] *Connectionist network models* (see chapter 4) are typically motivated by cognitive phenomena and are governed primarily by computational constraints, while figuring in the background are very general neurobiological constraints such as number of processing units and time required to perform a task. Accordingly they are more properly considered demonstrations than models, in the sense defined above. *Models of real neural networks,* by contrast, are primarily motivated by biological constraints, such as the physiological and anatomical properties of specific cell types (Arbib and George

1987, MacGregor 1987a,b, Sejnowski, Koch, and Churchland 1988, Koch and Segev 1989). Despite their different origins and sources of dominant constraints both connectionist models and neural models are based on the mathematics of nonlinear dynamical systems in high-dimensional spaces (Abraham and Shaw 1982). As a consequence, there is a common conceptual and technical background that allows researchers in both fields to communicate with each other. Such communication provides links between two rich sources of experimental data, and consequently connectionist and neural models have the potential to coevolve toward an integrated, coherent account of information processing in the mind-brain.

The ultimate goal of a unified account does not entail that it be a *single* model that spans all the levels of organization seen in nervous systems, or that the highest level be explained directly in terms of events at the molecular level. Instead it is more probable that the integration consist of a chain of theories and models that links adjacent levels. The unifying connections would derive, therefore, from the chain of interlocking theories in virtue of which phenomena at one level are explained in terms of phenomena at a lower level, and those in turn by phenomena at yet lower levels. Notice also that should one level be explained in terms of a lower level, this does not mean that the higher-level theory is no longer useful, or that the phenomena at that level no longer exist, or that if they do, they are no longer worth studying. On the contrary, such levels, and the theories pertinent to those levels, will persist. As in genetics and embryology, explanations coexist at all levels, from the molecular to the systems level.

8.5 Conclusions

It would be very convenient if we could understand the nature of cognition without understanding the nature of the brain itself. Unfortunately, it is very difficult if not impossible to theorize effectively on these matters in the absence of neurobiological constraints. The primary reason is that computational space is consummately vast, and there are many conceivable solutions to the problem of how a cognitive operation could be accomplished. Neurobiological data provide essential constraints on computational theories, and they consequently provide an efficient means for narrowing the search space. Equally important, the data are also richly suggestive of hints concerning what might really be going on and what computational strategies evolution might have chanced upon. Moreover, it is by no means clear or settled yet what exactly are the functional categories at the cognitive levels, and theories of lower-level function may well be crucial to the discovery of the nature of higher-level organization. Accordingly, despite the fact that the brain is experimentally demanding, neurobiology is indispensable to the discovery of theories that explain how we perform such tasks as seeing, thinking, and being aware.

On the other hand, the possibility that cognition will be an open book once we understand the details of each and every neuron and its development, connectivity, response properties, and so forth is likewise misconceived. Even if we could simulate, synapse for synapse, our entire nervous system, that accomplishment, by itself, would not be the same as understanding how it works. The simulation might be just as much of a mystery as the function of the brain currently is, for it may reveal nothing about the network and systems properties that hold the key to cognitive effects. Genuine theorizing about the nature of neurocomputation is therefore essential.

Assuming that there are a number of levels of organization in nervous systems, such that cognitive science specifically addresses higher levels whereas neuroscience typically addresses lower levels, we can acknowledge this joint effort by saying that the goal is to figure out how the *mind-brain* works. In this sense, an ultimate goal is the *reductive integration* of the psychological and neurobiological sciences, and thus cognitive neuroscience is a genuinely interdisciplinary undertaking (LeDoux and Hirst 1986). Reduction here does not entail elimination, any more than the reduction of chemistry to physics entails the elimination of chemical principles. On the contrary, an integrative reduction between theories at different levels can provide insights that enrich the principles at both levels (Churchland 1986).

Remarkable developments in the seventies and eighties in cognitive science, computational theory, and neuroscience have engendered a new, if cautious, optimism for achieving some measure of integration and explanatory unification of the various levels of organization in nervous systems. Until quite recently the immediate goals of neuroscientists and cognitive scientists were sufficiently distant from each other that their discoveries often seemed of merely academic significance to the other. The third element, computer science, was unpromising as a means of bridging the gap because the dominant model of computation, based on the Turing machine and the von Neumann architecture, did not relate at all to what was known about nervous systems at the level of signal processing. All this has begun to change quite dramatically, and there is gathering conviction among scientists that the time is especially propitious for a fruitful convergence of research from hitherto isolated fields. In light of these considerations, the prospects favor genuine progress in generating theories that honor neurobiological and psychological constraints to explain how networks of neurons achieve high-level effects.

The social and institutional background for studying the relationship between mind and brain is rapidly evolving. Cross-disciplinary contact and collaboration is increasing, there are meetings on such topics as computational neuroscience and biological cognition, new journals for cognitive neuroscience and neural computation are being founded, and departments of cognitive science and institutes for the study of the mind-brain are coming into existence. These developments signify the need for stable scientific institutions to enable further cross-disciplinary research, and they are enormously important if the integrative program is to succeed. For it is not enough that researchers in one field read the journals and attend the meetings of the other field. Rather it will be through working together on common projects that the major breakthroughs will most likely come.

Notes

We are indebted to Francis Crick, whose insights as well as critical judgments were a major resource in writing this chapter. We are grateful to Michael Posner and William Lytton for thoughtful readings of the manuscript.

1. The original concept of levels of analysis can be found in Marr and Poggio (1976, 1977). Although Marr (1982) emphasized the importance of the computational level, the notion of a hierarchy of levels grew out of earlier work by Reichardt and Poggio (1976) on the visual control of orientation in the fly. In a sense the current view on the interaction between levels is not so much a departure from the earlier views as a return to the practice that was previously established by Reichardt, Poggio, and even by Marr himself, who published a series of papers on neural-network models of the cerebellar cortex and cerebral cortex (see, for example, Marr 1969, 1970). The emphasis on the computational level has nonetheless had an important influence on the problems and issues that concern the current generation of neural and connectionist models (Sejnowski, Koch, and Churchland 1988).

2. We owe this analysis to Francis Crick, who sees demonstrations as "don't-worry theories," inasmuch as they address only the possibility that something could be done in a certain way, not the further question of whether it is in fact done in that way.

Selected Annotated Readings

Other chapters in this book touch on specialized aspects of brain and cognition including vision (chapter 15), attention (chapter 16), memory (chapter 17), and motor control (chapters 18 and 19). In this annotated bibliography we list some accessible books, journals, and reviews that can provide the reader with an entry into the literature. The list is selective and by no means exhaustive.

Arbib, M. A., 1987. *Brains, Machines and Mathematics.* 2nd ed. New York: Springer-Verlag. Historical survey of brain modeling.

Changeux, J.-P., 1985. *Neuronal Man.* Oxford: Oxford University Press. Sweeping view of the brain and behavior in the light of developmental and molecular neurobiology.

Churchland, Patricia Smith. 1986. *Neurophilosophy: Toward a Unified Science of the Mind-Brain.* Cambridge, MA: MIT Press. Exploration of the rapproachment of the cognitive and neural sciences.

Dowling, J. E. 1987. *The Retina: An Approachable Part of the Brain.* Cambridge, MA: Harvard

University Press. Though separate from the brain, the retina is part of the central nervous system, and it is one of the few parts some of whose functions we may already understand.

Groves, Philip M., and Rebec, George V. 1988. *Introduction to Biological Psychology.* 3rd ed. Dubuque, IA: William C. Brown Co. Publishers. Comprehensive and very readable basic introduction.

Hubel, D. H. 1988. *Eye, Vision and Brain.* New York: Freeman. Masterly survey of what is currently known about the structure and function of the visual systems of cats and monkeys.

Jeannerod, M. 1985. *The Brain Machine.* Cambridge, MA: Harvard University Press. Thoughtful historical account of the development of ideas on the neurophysiology of higher mental functions, particularly of action systems.

LeDoux, Joseph E., and Hirst, William, eds. 1986. *Mind and Brain: Dialogues in Cognitive Neuroscience.* Cambridge, Engl.: Cambridge University Press. Topics such as memory and attention are addressed by both psychologists and neuroscientists, with each commenting on the other. Illustrates the problems and prospects for interactions between neuroscientists and cognitive scientists.

Kandel, E., and Schwartz, J. 1985. *Principles of Neural Science.* 2nd ed. New York: Elsevier. Comprehensive introduction to all levels of brain organization including topics from neurology.

Kuffler, S. W., Nicolls, J. G., and Martin, A. R. 1984. *From Neuron to Brain: A Cellular Approach to the Function of the Nervous System.* 2nd ed. Sunderland, MA: Sinauer. Reliable, well-written introduction to cellular neurobiology.

Plum, F., and Mountcastle, V., eds. 1987. Higher functions of the brain. In *Handbook of Physiology.* Vol. 5, sec. 1. See especially chapters by R. Andersen on parietal cortex, P. Goldman-Rakic on frontal cortex, and S. Hillyard and T. Picton on evoked-responses potentials.

Rumelhart, D. G., and McClelland, J. L. 1986. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 1: Foundations. Volume 2: Psychological and Biological Models.* Cambridge, MA: MIT Press. Seminal collection of papers on connectionist modeling.

Schwartz, E., ed. 1990. *Computational Neuroscience.* Cambridge, MA: MIT Press. Collection of papers on modeling the nervous system, from the biophysical to the systems level.

Shepherd, G. M. 1987. *Neurobiology.* 2nd ed. Oxford: Oxford University Press. Broadly based survey of cellular and systems-level neuroscience, including invertebrate neurobiology.

Squire, L. R. 1987. *Memory and Brain.* Oxford: Oxford University Press. Accessible and sound entry to the literature on the biological basis of memory and learning.

Cognitive Neuroscience. Quarterly journal. Cambridge, MA: MIT Press. Research articles on neural mechanisms that underlie cognitive processing.

Neural Computation. Quarterly journal. Cambridge, MA: MIT Press. Short research papers

on brain models in computational neuroscience and neural information processing systems in neurocomputing. Also contains reviews that are accessible to a general reader.

Trends in Neuroscience. Monthly journal (Elsevier). Contains brief but very useful technical reviews and is a good place to get up-to-date references to the literature on topics of current interest.

Annual Reviews of Neuroscience. Palo Alto, CA: Annual Reviews Inc. Comprehensive reviews of the literature on special topics.

Daedalus: Journal of the American Academy of Arts and Sciences. Winter 1988 Special Issue on Artificial Intelligence. Includes many interesting essays on the relationship between brains and computers. See especially "Neural Nets and Artificial Intelligence," by J. D. Cowan and D. H. Sharp; "The New Connectionism: Developing Relationships Between Neuroscience and Artificial Intelligence," by J. T. Schwartz; "Real Brains and Artificial Intelligence," by G. N. Reeke Jr. and G. M. Edelman; "The Prospects for Building Truly Intelligent Machines," by D. L. Waltz; and "Making Machines (and Artificial Intelligence) See," by A. Hurlbert and T. Poggio.

References

Abraham, R. F., and Shaw, C. D. 1982. *Dynamics, the Geometry of Behavior.* Santa Cruz: Aerial Press.

Adrian, E. D. 1953. The mechanism of olfactory stimulation in the mammal. *Adv. Sci. (Lond.)* 9:417–420.

Alkon, D. L. 1987. *Memory Traces in the Brain.* Oxford: Oxford University Press.

Allman, J. 1982. Reconstructing the evolution of the brain in primates through the use of comparative neurophysiological and neuroanatomical data. In E. Armstrong and D. Falk, eds. *Primate Evolution.* New York: Plenum, pp. 13–28.

Allman, J., Miezin, F., and McGuiness, E. 1985. Stimulus specific response from beyond the classical receptive field: Neurophysiological mechanisms for local–global comparisons in visual neurons. In W. M. Cowan, E. M. Shooter, C. F. Stevens, and R. F. Thompson, eds. *Annual Review of Neuroscience.* Palo Alto: Annual Reviews, Inc., pp. 407–430.

Andersen, R. A., and Mountcastle, V. B. 1983. The influence of the angle of gaze upon the excitability of light-sensitive neurons of the posterior parietal cortex. *Journal of Neuroscience* 3:532–548.

Arbib, M., and George, S. A. 1987. *1987 Short Course Syllabus: Computational Neuroscience.* Washington, DC: Society for Neuroscience.

Asanuma, H. 1975. Recent developments in the study of the columnar arrangement of neurons within motor cortex. *Physiological Review* 55:143–156.

Bachus, R., Mueller, E., Koenig, H., Braeckle, G., and Weber, H. 1987. Functional imaging using NMR. In V. R. McCready, M. Leach, and P. J. Ell, eds. *Functional Studies Using NMR.* New York: Springer-Verlag, pp. 43–60.

Barlow, H. B. 1985. The Twelfth Bartlett Memorial Lecture: The role of single neurons in the psychology of perception. *Quart. J. Exp. Psych.* 37A: 121–145.

Berger, T. W., and Thompson, R. F. 1978. Neuronal plasticity in the limbic system during classical conditioning of the rabbit nictitating membrane response. 1. The hippocampus. *Brain Research* 145:323–346.

Blasdel, G. G., and Salama, G. 1986. Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature* 321:579–585.

Boltz, J., and Gilbert, C. D. 1986. Generation of end-inhibition in the visual cortex via interlaminar connections. *Nature* 320:362–365.

Boyde, A. 1985. Stereoscopic images in confocal (tandem scanning) microscopy. *Science* 230:1270–1272.

Brown, T. H., Ganong, A. H., Kariss, E. W., and Keenan, C. L. 1989. Hebbian synapses—Computations and biophysical mechanisms. *Ann. Rev. Neurosci.* 12: (in press).

Cajal, S. R. 1937. *Recollections of My Life.* Philadelphia: American Philosophical Society.

Campbell, C. B. G., and Hodos, W. 1970. The concept of homology and the evolution of the nervous system. *Brain, Behavior and Evolution* 3:353–367.

Churchland, P. M. 1985. Some reductive strategies in cognitive neurobiology. *Mind* 95:279–309.

Churchland, P. S. 1986. *Neurophilosophy: Toward a Unified Science of the Mind-Brain.* Cambridge, MA: MIT Press.

Churchland, P. S., and Sejnowski, T. J. 1988a. Neural representations and neural computations. In L. Nadel, ed. *Neural Connections and Mental Computation.* Cambridge, MA: MIT Press.

Churchland, P. S., and Sejnowski, T. J. 1988b. Perspectives in cognitive neuroscience. *Science* 242:741–745.

Connor, J. A., Tseng, H. S., and Hockberger, P. E. 1987. Depolarization- and transmitter-induced changes in intracellular calcium of rat cerebellar granule cells in explant cultures. *Journal of Neuroscience* 7:1384–1400.

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

Crick, F. H. C. 1979 (September). Thinking about the brain. *Scientific American* 241(3):219–232.

Crick, F. H. C., and Asanuma, C. 1986. Certain aspects of the anatomy and physiology of the cerebral cortex. In J. L. McClelland, and D. E. Rummelhart, eds. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Psychological and Biological.* Cambridge, MA: MIT Press, pp. 333–371.

Damasio, A. R. 1985. Disorders of complex visual processing: Agnosias, achromotopsia, Balint's syndrome, and related difficulties of orientation and construction. In M.-M. Mesulam, ed. *Principles of Behavioral Neurology.* Philadelphia: F.A. Davis, pp. 259–288.

DeYoe, E. A., and Van Essen, D. C. 1988. Concurrent Processing Streams on monkey visual cortex. *Trends in Neuroscience* 11:219–226.

Dobbins, A., Zucker, S. W., and Cynader, M. S. 1987. Endstopped neurons in the visual cortex as a substrate for calculating curvature. *Nature* 329:438–441.

Donchin, E., McCarthy, G., Kutas, M., and Ritter, W. 1983. Event-related potentials in the study of consciousness. In G. E. Schwartz and D. Shapiro, eds. *Consciousness and Self-Regulation.* New York: Plenum, pp. 81–121.

Feldman, J. A., and Ballard, D. H. 1982. Connectionist models and their properties. *Cognitive Science* 6:205–254.

Ferrier, D. (1876). *The Function of the Brain.* London: Smith, Elder.

Ferster, D., and Koch, C. 1987. Neuronal connections underlying orientation selectivity in cat visual cortex. *Trends in Neurosciences* 10:487–492.

Fetz, E. E., and Cheney, P. D. 1980. Postspike facilitation of forelimb muscle activity by primate corticomotoneural cells. *Journal of Neurophysiology* 44:751–772.

Fox, P. T., Fox, J. M., and Raichle, M. E. 1985. The role of the cerebral cortex in the generation of voluntary saccades: A positron emission tomographic study. *Journal of Neurophysiology* 54:348–369.

Fox, P. T., Miezin, F. M., Allman, J. A., Van Essen, D. C., and Raichle, M. E. 1987. Retinoptic organization of human visual cortex mapped with positron emission tomography. *Journal of Neuroscience* 7:913–922.

Gall, F. J., and Spurzheim, J. C. 1968. Anatomie et physiologie du system nerveux en general, et du cerveau en particulier, avec des observations sur la possibilité de reconnaître plusieurs dispositions intellectuelles et morales de l'homme et des animaux, par la configuration de leurs têtes. In E. Clark, and C. D. O'Mally, eds. *The Human Brain and Spinal Cord: A Historical Study Illustrated by Writings from Antiquity to the Twentieth Century.* Berkeley and Los Angeles: University of California Press, pp. 476–480.

Gerstein, G. L., and Aertsen, A. M. H. J. 1985. Representation of cooperative firing activity among simultaneously recorded neurons. *Journal of Neurophysiology* 54:1513–1528.

Gerstein, G. L., Bloom, M. J., Espinosa, I. E., Evanczuk, S., and Turner, M. R. 1983. Design of a laboratory for multineuron studies. *IEEE Transactions on Systems, Man, and Cybernetics* SMC-13:668–676.

Goldman-Rakic, P. S. 1987. Circuitry of primate prefrontal cortex and regulation of behavior by representational knowledge. In F. Plum, and V. Mountcastle, eds. *Higher Cortical Function: Handbook of Physiology.* Washington, DC: American Physiological Society, pp. 373–417.

Goldman-Rakic, P. S., and Selemon, L. D. 1986. Topography of corticostriatal projections in nonhuman primates and implications for functional parcellation of the neostriatum. In E. G. Jones, and A. Peters, eds. *Cerebral Cortex.* New York: Plenum Press, pp. 447–466.

Graybiel, A. N., and Hickey, T. L. 1982. Chemospecificity of ontogenetic units in the striatum: Demonstration by combining [^{3}H]thymidine neuronography and histochemical staining. *Proc. Natl. Acad. Sci. USA* 79:198–202.

Grinvald, A. 1985. Real-time optical mapping of neuronal activity: From single growth cones to the intact mammalian brain. In W. M. Cowan, E. M. Shooter, C. F. Stevens, and R. F. Thompson, eds. *Annual Review of Neuroscience.* Palo Alto: Annual Reviews, Inc., pp. 263–305.

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

Hanny, P. E., Maunsell, J. H. R., and Schiller, P. H. 1988. State-dependent activity in monkey visual cortex: II. Visual and nonvisual factors in V4. *Experimental Brain Research* 69:245–259.

Hillyard, S. A., and Picton, T. W. 1987. Electrophysiology of cognition. In F. Plum, ed. *Handbook of Physiology Section 1: Neurophysiology.* New York: American Physiological Society, pp. 519–584.

Hinton, G. E. 1986. Learning distributed representations of concepts. In *Proc. 8th Ann. Conf. Cog. Sci. Soc.* Hillsdale, NJ: Lawrence Erlbaum Assoc., pp. 1–12.

Hippocrates. 1949. *Ancient Medicine and Other Treatises.* F. Adams, trans. Chicago: Regnery.

Hobson, J. A. 1985. The neurobiology and pathophysiology of sleep and dreaming. *Discussions in Neuroscience* 2:9–50.

Hokfelt, T. 1987. Neuronal communication through multiple coexisting messengers. In G. M. Edelman, W. E. Gall, and W. M. Cowan, eds. *Synaptic Function.* New York: John Wiley & Sons, pp. 179–211.

Horton, J. C., and Sherk, H. 1984. Receptive field properties in the cat's lateral geniculate nucleus in the absence of on-center retinal input. *Journal of Neuroscience* 4:374–380.

Hubel, D. H., and Livingstone, M. S. 1987. Segregation of form, color, and stereopsis in primate area 18. *Journal of Neuroscience* 7:3378–3415.

Hubel, D. H., and Wiesel, T. N. 1962. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol. (Lond.)* 160:106–154.

Hubel, D. H., and Wiesel, T. N. 1977. Functional architecture of macaque monkey visual cortex. *Proc. R. Soc. Lond. B* 198:1–59.

Hurlbert, A., and Poggio, T. 1988. Synthesizing a color algorithm from examples. *Science* 239:482–485.

Ingvar, D. H., and Schwartz, M. S. 1974. Blood flow patterns induced in the dominant hemisphere by speech and reading. *Brain* 97:273–288.

Jackson, J. H. 1958. *Selected Writings of John Hughlings Jackson.* New York: Basic Books.

Jones, E. G., Juliano, S. L., and Whitsel, B. L. 1987. A combined 2-deoxyglucose and neurophysiological study of primate somatosensory cortex. *Journal of Comparative Neurology* 263:514–525.

Kaas, J. H., Nelson, R. J., Sur, M., and Merzenich, M. M. 1979. Multiple representations of the body within the primary somatosensory cortex of primates. *Science* 204:521–523.

Kandel, E., and Schwartz, J. 1985. *Principles of Neural Science.* 2nd ed. New York: Elsevier.

Kandel, E. R. 1983. Neurobiology and molecular biology: The second encounter. In *Cold Spring Harbor Symposia on Quantitative Biology.* Cold Spring Harbor: Cold Spring Harbor Laboratory, pp. 891–908.

Kandel, E. R., Klein, M., Hochner, B., Schuster, M., Siegelbaum, S. A., Hawkins, R. D., Glanzman, D. L., and Castellucci, V. F. 1987. Synaptic modulation and learning. New insights into synaptic transmission from the study of behavior. In G. M. Edelman, W. E. Gall, and W. M. Cowan, eds. *Synaptic Function.* New York: John Wiley and Sons, pp. 471–518.

Knudsen, E. I., du Lac, S., and Esterly, S. D. 1987. Computational maps in the brain. In W. M. Cowan, E. M. Shooter, C. F. Stevens, and R. F. Thompson, eds. *Annual Review of Neuroscience.* Palo Alto: Annual Reviews, Inc., pp. 41–65.

Koch, C., and Poggio, T. 1987. Biophysics of computation: Neurons, synapses, and membranes. In G. M. Edelman, W. E. Gall, and W. M. Cowan, eds. *Synaptic Function.* New York: John Wiley & Sons, pp. 637–697.

Koch, C., and Segev, I. 1989. *Methods in Neuronal Modeling: From Synapse to Networks.* Cambridge, MA: MIT Press.

Konishi, M. 1986. Centrally synthesized maps of sensory space. *Trends in Neurosciences* 9:163–168.

Kosslyn, S. M. 1987. Seeing and imaging in the cerebral hemispheres: A computational approach. *Psych. Rev.* 94:148–175.

Kosslyn, S. M. 1988. Aspects of cognitive neuroscience of mental imagery. *Science* 240:1621–1626.

Kosslyn, S. M., Holtzman, J. D., Gazzaniga, M. S., and Farrah, M. J. 1985. A computational analysis of mental imagery generation: Evidence for functional dissociation in split brain patients. *J. Exp. Psych.: General* 114:311–341.

Kuffler, S. W. 1980. Slow synaptic responses in autonomic ganglia and the pursuit of a peptidergic transmitter. In: E. A. Kravitz and J. E. Treherne, eds. *Neurotransmission, Neurotransmitters, and Neuromodulators.* Cambridge, Engl.: Cambridge University Press, pp. 257–286.

Kutas, M., and Van Petten, C. 1988. Event-related brain potential studies of language. In P. K. Ackles, J. R. Jennings, and M. G. H. Coles, eds. *Advances in Psychophysiology.* Greenwich, CT: JAI Press.

Land, E. H., Hubel, D. H., Livingstone, M. S., Perry, S. H., and Burns, M. M. 1983. Colour-generating interactions across the corpus callosum. *Nature* 303:616–618.

LeDoux, J., and Hirst, W. 1986. *Mind and Brain: Dialogues in Cognitive Neuroscience.* Cambridge, Engl.: Cambridge University Press.

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

Lisberger, S. G., Morris, E. J., and Tychsen, L. 1987. *Ann. Rev. Neurosci.* 10:97–129.

Livingstone, M. S., and Hubel, D. H. 1987a. Connections between layer 4B of area 17 and the thick cytochrome oxidase stripes of area 18 in the squirrel monkey. *Journal of Neuroscience* 7:3371–3377.

Livingstone, M. S., and Hubel, D. H. 1987b. Psychophysical evidence for separate channels for the perception of form, color, movement, and depth. *Journal of Neuroscience* 7:3416–3468.

Llinas, R. R. 1985. Electronic transmission in the mammalian central nervous system. In M. E. L. Bennett, and D. C. Spray, eds. *Gap Junctions.* Cold Spring Harbor: Cold Spring Harbor Laboratory, pp. 337–353.

Lund, J. S. 1987. Local circuit neurons of macaque monkey striate cortex: I. Neurons of laminae 4C and 5A. *Journal of Comparative Neurology* 257:60–92.

Luyten, P. R., and den Hollander, J. A. 1986. Observation of metabolites in the human brain by MR spectroscopy. *Radiology* 161:795–798.

MacGregor, R. J. 1987a. Simplified models of single neurons. In *Neural and Brain Modeling.* New York: Harcourt Brace Jovanovich, pp. 220–260.

MacGregor, R. J. 1987b. Synaptic bombardment in model neurons. In *Neural and Brain Modeling.* New York: Harcourt Brace Jovanovich, pp. 261–288.

Mandelbrot, B. B. 1983. *The Fractal Geometry of Nature.* San Francisco, W. H. Freeman.

Marr, D. 1969. A theory of cerebellar cortex. *J. Physiol. (Lond.)* 202:437–470.

Marr, D. 1970. A theory for cerebral neocortex. *Proc. R. Soc. Lond. B* 176:161–234.

Marr, D. 1982. Vision. San Francisco: W. H. Freeman.

Marr, D., and Poggio, T. 1976. *From Understanding Computation to Understanding Neural Circuitry.* MIT Artificial Intelligence Laboratory Technical report. AI Memo 357, MIT AI Laboratory, Cambridge, MA.

Marr, D., and Poggio, T. 1977. From understanding computation to understanding neural circuitry. *Neurosciences Res. Prog. Bull.* 15:470–488.

Martin, K. A. C. 1984. Neuronal circuits in cat striate cortex. In E. G. Jones, and A. Peters, eds. *Cerebral Cortex.* New York: Plenum Press.

Maunsell, J. H. R., and Newsome, W. T. 1987. Visual processing in monkey extrastriate cortex. In W. M. Cowan, E. M. Shooter, C. F. Stevens, and R. F. Thompson, eds. *Ann. Rev. Neurosci.* Palo Alto: Annual Reviews Inc., pp. 363–401.

Maunsell, J. H. R., and Van Essen, D. C. 1983. The connections of the middle temporal visual area (MT) and their relationship to a cortical hierarchy in macaque monkey. *Journal of Neuroscience* 3:2563–2586.

McCarthy, R. A., and Warrington, E. K. 1988. Evidence for modality-specific meaning systems in the brain. *Nature* 334:428–430.

McNaughton, B. L., and Morris, R. G. 1987. Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neurosciences* 10:408–415.

Merzenich, M. M., and Brugge, J. F. 1973. Representation of the cochlear partition on the superior temporal plane in the macaque monkey. *Brain Research* 50:275–296.

Merzenich, M. M., and Kaas, J. H. 1982. Reorganization of mammalian somatosensory cortex following peripheral nerve injury. *Trends in Neurosciences* 5:434–436.

Mesulam, M. M. 1985. Attention, confusional states, and neglect. In M. M. Mesulam, ed. *Principles of Behavioral Neurology.* Philadelphia: F. A. Davis, pp. 125–168.

Milner, B. 1966. Amnesia following operation on the temporal lobes. In C. W. M. Whitty and O. Zangwill, ed. *Amnesia.* London: Butterworth, pp. 109–133.

Mishkin, M. 1982. A memory system in the monkey. *Phil. Trans. R. Soc. Lond. B* 298:85–95.

Moran, J., and Desimone, R. 1985. Selective attention gates visual processing in the extrastriate cortex. *Science* 229:782–784.

Newsome, W. T., and Pare, E. B. 1986. MT lesions impair visual discrimination of direction in a stochastic motion display. *Society of Neuroscience Abstracts* 12:1183.

Newsome, W. T., Wurtz, R. H., Durtsteler, M. R., and Mikami, A. 1985. Deficits in visual motion processing following ibotenic acid lesions of the middle temporal visual area of the macaque monkey. *J. Neurosci.* 5:825–840.

O'Keefe, J., and Nadel, L. 1978. *The Hippocampus as a Cognitive Map.* Oxford: Clarendon Press.

Paller, K. A., Zola-Morgan, S., Squire, L. R., and Hillyard, S. A. 1984. Monkeys with lesions of hippocampus and amygdala exhibit event-related brain potentials that resemble the human P300 view. *Soc. Neurosci. Abstr.* 10:849.

Paller, K. A., Zola-Morgan, S., Squire, L. R., and Hillyard, S. A. 1989. P-3 like brain waves in normal monkeys and monkeys with medial temporal lesions. *Behavioral Neuroscience* 102:714–725.

Penfield, W., and Jasper, H. 1954. *Epilepsy and the Functional Anatomy of the Human Brain.* Boston: Little, Brown.

Pentland, A. P. 1984. Fractal-based description of natural scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI-6:661–674.

Perrett, D. I., Mistlin, A. J., and Chitty, A. J. 1987. Visual neurones responsive to faces. *Trends in Neurosciences* 10:358–364.

Petersen, S. E., Fox, P. T., Posner, M. I., Mintun, M. A., and Raichle, M. E. 1988. Positron emission tomographic studies of the cortical anatomy of single-word processing. *Nature* 331:585–589.

Phelps, M. E., and Mazziotta, J. C. 1985. Positron emission tomography: Human brain function and biochemistry. *Science* 228:799–809.

Posner, M. I., Petersen, S. E., Fox, P. T., and Raichle, M. E. 1988. Localization of cognitive operations in the human brain. *Science* 240:1627–1631.

Purves, D., and Voyvodic, J. T. 1987. Imaging mammalian nerve cells and their connections over time in living animals. *Trends in Neurosciences* 10:398–404.

Raichle, M. E. 1986. Neuroimaging. *Trends in Neurosciences* 9:525–529.

Reichardt, W., and Poggio, T. 1976. Visual control of orientation behavior in the fly. *Quarterly Reviews of Biophysics* 9:311–375.

Reitboeck, H. J. P. 1983. A 19-channel matrix drive with individually controllable fiber microelectrodes for neurophysiological applications. *IEEE Transactions on Systems, Man, and Cybernetics* SMC-13:676–683.

Reiter, H. O., and Stryker, M. P. 1987. A novel expression of plasticity in kitten visual cortex in the absence of postsynaptic activity. *Society for Neuroscience Abstracts* 13:1241.

Ritchie, J. M. 1987. Ion channels in neural membranes. *Discussions in Neurosciences* 4:11–55.

Robinson, D. A. 1981. The use of control systems analysis in the neurophysiology of eye movements. *Annual Review of Neuroscience* 4:463–503.

Roland, P. E. 1984a. Organization of motor control by the normal human brain. *Human Neurobiology* 2:205–216.

Roland, P. E. 1984b. Somatotopic tuning of postcentral gyrus during focal attention in man. *Journal of Neurophysiology* 46:744–754.

Salzberg, B. M., Obaid, A. L., Senseman, D. M., and Gainer, H. 1983. Optical recording of action potentials from vertebrate nerve terminals using potentiometric probes provides evidence for sodium and calcium components. *Nature* 306:36–40.

Scherg, M., and Von Cramon, D. 1986. Evoked dipole source potentials of the human auditory cortex. *Electroenceph. Clin. Neurophys.* 65:344–360.

Schiller, P. 1982. The central connections of the retinal on and off pathways. *Nature* 297:580–583.

Sejnowski, T. J. 1986. Open questions about computation in cerebral cortex. In: J. L. McClelland and D. E. Rumelhart, eds. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition.* Vol. 2. Cambridge, MA: MIT Press, pp. 372–389.

Sejnowski, T. J. 1987. Computational models and the development of topographic projections. *Trends in Neurosciences* 10:304–305.

Sejnowski, T. J. 1988. Neural network learning algorithms. In R. Eckmiller, and C. von der Malsberg, eds. *Neural Computers.* New York: Springer-Verlag.

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

Sejnowski, T. J., and Tesauro, G. J. 1988. The Hebb rule for synaptic plasticity: Implementations and applications. In J. Byrne and W. O. Berry, eds. *Neural Models of Plasticity.* New York: Academic Press.

Sereno, M. I. 1988. The visual system. In I. W. V. Seelen, U. M. Leinhos, G. Shaw, eds. *Organization of Structure and Function in the Brain.* Basel: VCH Verlagsgesellschaft.

Shepherd, G. M. 1988. The basic circuit of cortical organization. In M. S. Gazzaniga, ed. *Perspectives in Memory Research.* Cambridge, MA: MIT Press.

Shepherd, G. M. 1989. The significance of real neuron architectures for neural network simulations. In E. Schwartz, ed. *Computational Neuroscience.* Cambridge, MA: MIT Press.

Shepherd, G. M., Brayton, R. K., Miller, J. P., Segev, I., Rinzel, J., and Rall, W. 1985. Signal enhancement in distal cortical dendrites by means of interactions between active dendritic spines. *Proceedings of the National Academy of Sciences USA* 82:2192–2195.

Sherrington, C. 1940. *Man and his Nature.* Cambridge, Engl.: Cambridge University Press.

Sherrington, C. S. 1906. *The Integrative Action of the Nervous System.* New Haven: Yale University Press.

Siegel, R. M., and Andersen, R. A. 1986. Perceptual deficits following ibotenic acid lesions of the middle temporal area (MT) in the behaving rhesus monkey. *Society for Neuroscience Abstracts* 12:1183.

Sokoloff, L. 1984. *Metabolic Probes of Central Nervous System Activity in Experimental Animals and Man.* Sunderland, MA: Sinauer Associates.

Sperry, R. W., and Gazzaniga, M. 1967. Language following surgical disconnection of the hemispheres. In C. Millikan, and F. Darley, eds. *Brain Mechanisms Underlying Speech and Language.* New York: Grune and Stratton, pp. 108–115.

Squire, L. R. 1987. *Memory and Brain.* Oxford: Oxford University Press.

Squire, L. R., Shimamura, A. P., and Amaral, D. G. 1988. Memory and the hippocampus. In J. Byrne, and W. Berry, eds. *Neural Models of Plasticity.* New York: Academic Press.

Sterling, P. 1983. Microcircuitry of the cat retina. In W. M. Cowan, E. M. Shooter, C. F. Stevens, and R. F. Thompson, eds. *Annual Review of Neuroscience.* Palo Alto: Annual Reviews, Inc., pp. 149–185.

Stewart, W. B., Kauer, J. S., and Shepherd, G. M. 1979. Functional organization of rat olfactory bulb analyzed by the 2-deoxyglucose method. *Journal of Comparative Neurology* 185:715–734.

Suga, N. 1984. The extent to which biosonar information is represented in the bat auditory cortex. In G. M. Edelman, W. E. Gall, and W. M. Cowan, eds. *Dynamic Aspects of Neocortical Function*. New York: John Wiley & Sons, pp. 315–373.

Swindale, N. V., Matsubara, J. A., and Cynader, M. S. 1987. Surface organization of orientation and direction selectivity in cat area 18. *Journal of Neuroscience* 7:1414–1427.

Tootell, R. B. H., Silverman, M. S., Switkes, E., and De Valois, R. L. 1982. Deoxyglucose analysis of retinotopic organization in primate striate cortex. *Science* 218:902–904.

Ts'o, D. Y., Gilbert, C. D., and Wiesel, T. N. 1986. Relationship between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *Journal of Neuroscience* 6:1160–1170.

Tsien, R. Y., and Poenie, M. 1986. Fluorescence ratio imaging: A new window into intracellular ionic signaling. *Trends in Biochemical Sciences* 11:450–455.

Van Petten, C., and Kutas, M. 1987. Ambiguous words in context. An event-related potential analysis of the time course of meaning activation. *Journal of Memory and Language* 26:188–208.

von Helmholtz, H. 1948. On the rate of transmission of the nerve impulse. In W. Dennis, ed. *Readings in the History of Psychology*. New York: Appleton-Century-Crofts, pp. 197–198.

von der Heydt, R., Peterhans, E., and Baumgartner, G. 1984. Illusory contours and cortical neuron responses. *Science* 224:1260–1262.

Williamson, S. J., and Kaufman, L. 1987. Analysis of neuromagnetic signals. In: A. Gevins and A. Rémond, eds. *Handbook of Electroencephalography and Clinical Neurophysiology*. Amsterdam: Elsevier.

Woolsey, T. A., and Van der Loos, H. 1970. The structural organization of layer IV in somatosensory region (SI) of mouse cerebral cortex. *Brain Research* 17:205–242.

Zeki, S. 1983. Colour coding in the cerebral cortex: The reaction of cells in monkey visual cortex to wavelengths and colours. *Neuroscience* 9:741–765.

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

Zola-Morgan, S., and Squire, L. R. 1984. Preserved learning in monkeys with medial temporal lesions. *Journal of Neuroscience* 4:1072–1085.