PART V

BIOLOGICAL MECHANISMS

One of the appeals of parallel distributed processing is the fact that it seems closer to the neural basis of cognition than most other approaches to cognitive processes. The idea that intelligent processing can emerge from the interactions of a large number of simple computational units and their interactions is, of course, directly inspired by what we know about the way the brain works. Further, several of the specific assumptions embodied in particular models have been inspired by observed characteristics of neural function.

This part of the book explores the neural mechanisms underlying parallel distributed processing in three different, though interrelated ways. First, it examines what we know and do not know about relevant aspects of the brain. Second, it considers the precise nature of the relation between concepts at a neural level of analysis, and concepts at higher, more cognitive levels. Third, it presents three specific attempts to bring PDP and neuroscience closer together by capturing aspects of neural function in simulation models based on the PDP framework laid out in Chapter 2.

The relevant neurophysiology. There is, of course, a vast body of data about the details of human brain function, and it would be hopeless to try to summarize all of this information within the confines of our book. However, among these data are some emerging principles and observations about the characteristics of the mammalian brain that seem particularly relevant to parallel distributed processing.

In Chapter 20, Crick and Asanuma describe the principles and observations they see as most salient, based on their ongoing survey of the neurophysiology of the mammalian brain. The focus is on the details of neuronal interaction and on the developing state of knoweldge about the regional architecture of the visual system, as revealed through physiological and anatomical experiments, primarily in primates. The material in this chapter is useful background for any PDP modeler, and will be of particular interest to anyone interested in modeling the detailed properties of real neural circuitry.

In Chapter 21, Sejnowski takes a somewhat different approach, and considers several general questions about brain function. By and large, neurophysiological data do not answer such questions in a definitive way, but they do suggest several principles of brain function that should continue to help steer our efforts to capture the essential features of the computational processes that take place in the human brain.

The relation between PDP models and the brain. There are several different ways in which PDP models relate to the brain. Different pieces of work represented in this book have different relationships to the details of brain structure. All share in common that they are "neurally inspired" models and are explorations of "brain-style" processing. Beyond this, however, there are some important differences. Here, we briefly outline some of these different approaches, and say how they relate to the chapters in Part V.

Some PDP models are intended to explore the computational properties of "brain-like" networks. The models described in Parts II and III are generally of this type. They consider sets of units with specified characteristics, and study what can be computed with them, how many units are required to do certain computations, etc. These models focus on parallel processing mechanisms more or less in their own right, quite apart from facts about the details of brain structure or the details of human behavior and cognition.

There are two principle kinds of motivations for this computational type of work. One is that the mechanisms under study are sufficiently brain-like that they may shed light on the actual way in which computation is done in the brain. The fact that these models generally idealize various properties of the brain may in many cases be a virtue for understanding brain function, since idealization, as Sejnowski points out in Chapter 21, can often facilitate an analysis that leads to deeper understanding of the emergent properties of complex systems. Indeed, some of this work is driven by the goal of exploring the implications of specific properties of brain function, such as the stochastic nature of neural firing, as explored in harmony theory (Chapter 6) and in Boltzmann machines (Chapter 7). The other motivation is to explore

the computational capacities of networks that appear on their surface to be well suited to certain kinds of information-processing tasks, such as search or representation building.

Other PDP models—and here we have in mind primarily the models in Part IV—attempt more directly to build accounts of human information processing capabilities, at a level of abstraction somewhat higher than the level of individual neurons and connections. In such models, the relationship between particular brain structures and particular elements of the models is not generally specified, but since the models attempt to capture the behavioral products of the activity of human brains, it is assumed that there is some relationship to real activity in the brain. The basic idea is that there is a mapping between elements of the model and the brain, but it is unknown and probably only approximate. A single unit may correspond to a neuron, a cluster of neurons, or a conceptual entity related in a complex way to actual neurons.

It might be thought that by adopting this more abstract approach, these models lose all contact with the underlying neurophysiological mechanisms. This is not the case. While models of cognitive processes may be developed without detailed regard for the underlying physiology, some of the characteristics of the brain clearly place constraints on the cognitive mechanisms. Some examples are the speed and precision of the basic computational elements, the general characteristics of their patterns of interconnection, the nature of the operations they can perform, the number of elements available in the brain, etc. (Chapter 4 provides a fuller discussion of these points). It becomes important, then, to develop some way of relating the more abstract, cognitive-level theory to the underlying neurophysiology. More fundamentally, this relation is central to conceptions of the relation between mind an brain. It is therefore of considerable importance to have an explicit theoretical framework for conceptualizing the exact nature of this relation. This point is addressed by Smolensky in Chapter 22.

In still other cases, the goal is to do neural modeling—to account for the facts of neuroscience rather than (or in addition to) the facts of cognition. This use of PDP models involves less idealization of the neural elements and more attention to the details of brain structure. Crick favors these applications for he feels that building up from the facts about real neural function is the best way to find out the way things really work—as opposed to the way things might work—in the brain. The last three chapters of Part V take this approach.

We hasten to add that we do not think any one of these uses is the "right" or "only" way to employ PDP models. Rather, we believe that work at each of these levels complements and reinforces work at the other levels, and that work at all of these levels will eventually allow us

to converge on an understanding of the nature of the cognitive processes that the brain supports and the neural mechanisms underlying these processes. We believe that the use of the common PDP framework for all of these applications will facilitate this process.

PDP models of neural mechanisms. Chapters 23 through 25 each take their own approach to modeling information-processing activity in the brain. Chapters 23 and 24 are explicit attempts to develop models that capture aspects of what is known about the behavior of neurons, while Chapter 25 focuses on neuropsychological data.

In Chapter 23, Zipser takes as his goal the development of biologically plausible models of place learning and goal localization in the rat. The goal of the models is to account for localization and place learning behavior, and at the same time, incorporate knowledge gained from single-unit recording experiments. The first model described in the chapter accounts for the behavior of so-called "place-field" cells in the hippocampus of the rat, and is closely tied to the physiological data. Two other, more speculative models work toward an account of goal localization, about which less of the physiology is known.

In Chapter 24, Munro considers the plasticity of neural mechanisms, as revealed through studies of single units in visual cortex after various schedules of visual deprivation and other forms of intervention. He shows how a very simple neural model that focuses on the plasticity of individual units can account for much of the data on the critical period. The account is based on the simple observation that changes in the connections of a neuron will make more difference if its prior connections are weak than if they are strong. The critical period is seen, on this account, as a simple manifestation of the natural consequences of the strengthening of connections through experience, and not as a manifestation of some sort of preordained maturational process that turns off plasticity.

Chapter 25 is not concerned with modeling data on the behavior of individual neurons; rather, it is concerned with reconciling neuropsychological evidence about amnesia with distributed models of memory. In distributed models, such as the one described in Chapter 17, information of different ages is stored in superimposed form, in the same set of connections. This fact provides a natural way of accounting for one aspect of amnesia: the fact that amnesics exhibit the residual ability to learn, gradually, from repeated experiences, even though their memory for individual episodes is extremely weak. Distributed memory, however, seems incompatible with another aspect of amnesia: namely, the temporally graded nature of the retrograde amnesia—the loss of prior information—that accompanies the reduction in the capacity to learn new material. If all memories are stored in the same set of

connections, why should more recent ones be more susceptible to loss or disruption than older ones? Chapter 25 reports simulations of various aspects of anterograde and retrograde amnesia, based on one possible answer to this question.

CHAPTER 20

Certain Aspects of the Anatomy and Physiology of the Cerebral Cortex

F. CRICK and C. ASANUMA

Our aim in this chapter is to describe some aspects of our present knowledge of the anatomy and physiology of the cerebral cortex of higher animals which may be of interest to theorists. We shall assume that readers have at least an elementary knowledge of this subject, so that they know, for instance, about the structure of neurons and the basis of neuronal excitation and synaptic transmission. The text by Kandel and Schwartz (1981) could be used to cover this background knowledge.

It is clearly impossible to describe most of what is known, even though this represents a tiny fraction of what one would like to know. We shall select examples to illustrate the general points we want to make. It will soon emerge that while some things are known with reasonable certainty, much is unknown or, even worse, surmised only on rather incomplete evidence. For this reason alone, the object of this chapter is not to dictate to theorists what "units" they must use in their modeling. It might turn out that theory will show that a particular process, or implementation of a process, gives a very advantageous performance, even though the experimentalists can, as yet, see no sign of it. The wise thing at that point would be to look for it experimentally, since it may have been overlooked for one or another technical reason. This aside, theorists should at least try to learn whether the features they wish to use for their implementation do actually occur in the relevant part of the brain, and they should be duly cautious if the

experimentalist can see no trace of them. Whether a theorist's unit can be a group of neurons is discussed later.

One other general point should perhaps be stated at the outset. Different parts of the brain are "wired" in radically different ways. It is thus not sensible to take one feature from, say, the olfactory bulb, another from the thalamus, and a third from the cerebellum, and combine them all together to account for a task that the cortex is expected to perform. Wherever possible, therefore, we shall choose examples from the mammalian cerebral cortex, both because so much work has been done on it and also because the problems theorists choose are often taken from aspects of human behavior that are mediated by the cerebral cortex. Excursions to other parts of the nervous system, such as the retina, cerebellum, and the olfactory bulb, will be made only when necessary to clarify certain points. Figure 1 illustrates a human brain and demonstrates the general location of some of its internal structures in relation to the cerebral cortex. However our aim in this chapter is not to describe the cerebral cortex as fully as possible, as one would need to do if one were concerned with its detailed workings, but merely to point out certain features of the cortex which should not be overlooked in theoretical modeling.

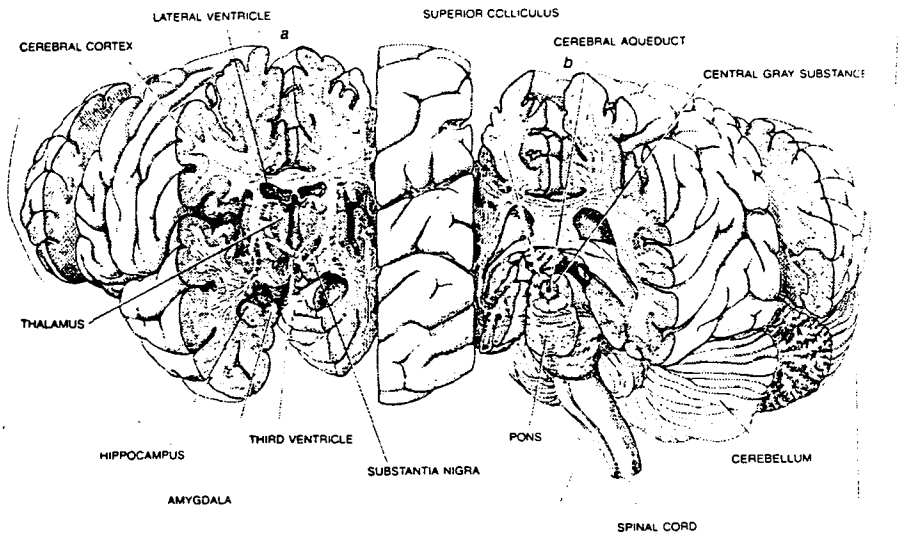

FIGURE 1. The human brain. The cerebral cortex is depicted transparently in this drawing so that some of the internal brain structures are visible. (From "The Organization of the Brain" by W. J. H. Nauta and M. Feirtag, 1979, *Scientific American*, *241*, p. 102. Copyright 1979 by W. H. Freeman & Co. Reprinted by permission.)

Another general point that should be made is that in many cases theorists choose problems associated with either language or the human visual system without recognizing that there is at least one important difference between them. Put briefly, there is no animal model for language, nor is it possible to carry out many types of experiments on the language centers of the human brain for obvious ethical reasons. Most of the really useful new methods used in neuroanatomy, such as tritiated amino acid autoradiography, horseradish peroxidase histochemistry, and metabolic mapping with [^{14}C] deoxyglucose can only be used effectively on animals. We are in the embarrassing position of knowing a lot about the neuroanatomy of the macaque monkey while having only a very limited amount of similar information about the human brain. Similarly, the most powerful neurophysiological technique—the use of microelectrodes for isolating the electrical activity of single neurons (or small groups of neurons)—is not suited for extensive use on humans. This disadvantage is partly offset by the greater ease with which human psychophysical experiments can be done. There are also a number of techniques which can be used to study aspects of the neural activity from the outside. These include position emission tomography (PET scanning), magnetic field detectors, electroencephalography, (EEG) and scalp recordings of evoked potentials. Unfortunately either the spatial or the temporal resolution of these methods is usually inadequate, and, as a result, the interpretation of the results is often not clear cut.

In the long run, a theoretical model in biology can only be validated by a *detailed* comparison with experiment. All psychophysical tests show that the performance of the visual system of the macaque monkey is roughly comparable to our own. From this point of view, therefore, the solutions of visual problems should be easier to bring down to earth than linguistic ones. This does not mean that linguistic problems may not suggest valuable ideas about the working of the brain. It does mean that they may be more difficult to test at the level of neuronal organization and function.

The Neuron

The "classical" neuron has several dendrites, usually branched, which receive information from other neurons and a single axon which outputs the processed information usually by the propagation of a "spike" or an "action potential." The axon ramifies into various branches that make synapses onto the dendrites and cell bodies of other neurons.

This simple picture (Figure 2A) has become complicated in several ways: (For a more thorough, yet general account, see the book by Shepherd, 1979.)

- A neuron may have no obvious axon but only "processes" that seem to both receive and transmit information (Figure 2B). An example of such neurons is the various amacrine cells found in the retina (Cajal, 1892). Although neurons without axons also occur in the olfactory bulb (Cajal, 1911), they have not been convincingly demonstrated in other parts of the nervous system.

- Axons may form synapses on other axons. In the cerebral cortex these synapses have been found only upon the *initial* segments of the axons of certain cells (Figure 2C) (Peters, Proskauer, & Kaiserman-Abramof, 1968; Westrum, 1966).

- Dendrites may form synapses onto other dendrites (Figure 2D). Examples of this are known in the retina (Dowling & Boycott, 1966), the olfactory bulb (Rall, Shepherd, Reese, & Brightman, 1966), the thalamus (Famiglietti, 1970), the superior colliculus (R. D. Lund, 1972), and the spinal cord (Ralston, 1968), but such contacts appear to be rare or absent in the cerebral cortex.

- An axon may not propagate a spike but instead produce a graded potential. Because of attenuation, we should expect this form of information signaling not to occur over long distances, and indeed it is found largely in such places as the retina, where the distances between connected neurons are shorter than in many other neural tissues; possibly because the time requirements are different (Figure 2E) (Werblin & Dowling, 1969). It is also conceivable that graded potentials occur at more local levels (Figure 2F). For example, an axon terminal forming a synapse on a given cell may itself receive a synapse. The presynaptic synapse may exert only a local potential change which is therefore restricted to that axon terminal. (The existence of this sort of a mechanism has been suggested for the spinal cord [Kuno, 1964] and the thalamus [Andersen, Brooks, Eccles, & Sears, 1964], but to date, no examples of this arrangement have been reported in the cerebral cortex.)

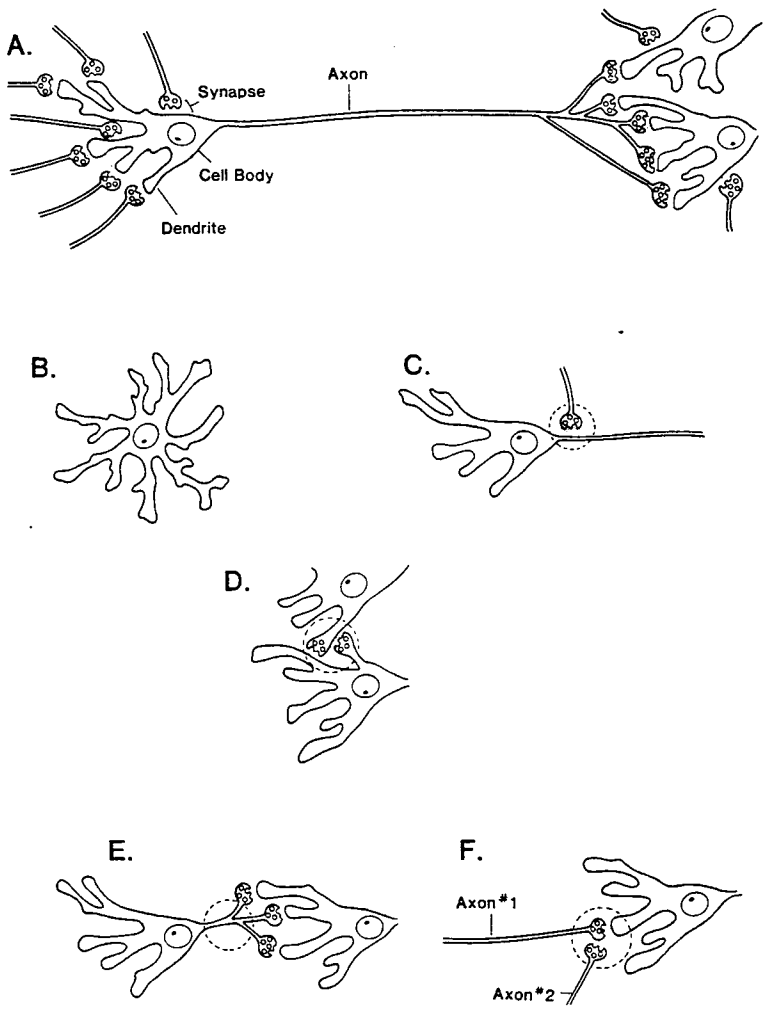

FIGURE 2. Highly schematized diagrams of the "classical" neuronal profile (*A*) and some of its variants (*B-F*). *A*: The "classical" neuron receives synapses on its dendrites and generates action potentials which travel down the axon. The axon subsequently branches and forms synapses on the dendrites and cell bodies of other neurons. *B*: There are neurons with no obvious axons. *C*: The initial segments of axons of neurons in the cerebral cortex may receive synapses. Note the highly strategic position of this kind of synapse. *D*: Dendrites forming synapses directly onto the dendrites of other neurons occur in the olfactory bulb and the thalamus. *E*: Graded potentials (instead of action potentials) can be effective if the axon is short. *F*: Graded potentials can also be effective at local levels. Here, Axon #2 can modulate the efficacy of the synapse formed by Axon #1 by producing a local potential change in the terminal of Axon #1.

Synapses

The great majority of synapses in the cerebral cortex are chemical, not electrical. A small star-shaped neuron (stellate cell) may receive a few hundred synapses, a small pyramid-shaped neuron (pyramidal cell) some thousands, and a large pyramidal cell some tens of thousands of synapses. Despite the large and variable number of synaptic contacts present upon neurons in the cerebral cortex, most synaptic contacts can be classified morphologically into two basic types (see Figure 3) (Peters, Palay, & Webster, 1976):

- *Type I.* These synapses have asymmetrical membrane specializations (the membrane thickening is greater on the postsynaptic side), and the presynaptic process contains fairly large (ca. 50 nm), round synaptic vesicles—believed to contain quanta, or packets of neurotransmitter. The synaptic cleft is usually about 30 nm across.

- *Type II.* These have symmetrical membrane specializations. The synaptic vesicles are smaller and, with the usual fixatives used for electron microscopy, are often ellipsoidal or flattened. (The shape of the vesicles depends on the details of the fixation and is not always a completely reliable criterion when comparing results reported by different investigators.) The synaptic cleft is usually 20 nm across and the zone of apposition is usually smaller than that of the Type I synapse.

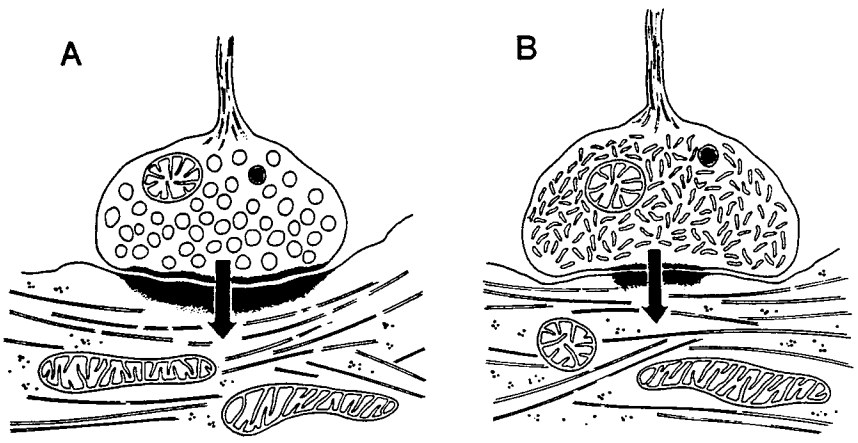

FIGURE 3. Idealized diagrams of a Type I (*A*) and a Type II (*B*) synapse. See text for explanation.

The importance of the classification into the two morphological types (originally recognized by Gray, 1959) is that *Type I synapses seem to be excitatory, whereas Type II synapses seem to be inhibitory.* We should add that the foregoing statement, though generally accepted by experimentalists, has not been systematically tested. In systems where the excitatory or inhibitory nature of a given synapse is well established, this correlation of morphology with physiology appears to be absolute (Uchizono, 1965).

There is another possible criterion for determining the character of synapses: This is the transmitter they use. In general, one is apt to assume that a given transmitter (or apparent transmitter) will usually do the same thing in different places, though there are well-established exceptions (depending on the nature of the postsynaptic receptors). Glutamate and aspartate always seem to excite, GABA (gamma-amino butyric acid) always seems to inhibit (Krnjević & Phillis, 1963). (It may come as a surprise to the reader to learn that for most cells in the brain we do not yet know what neurotransmitter they use.) The identity of the transmitters is usually determined immunocytochemically. Thus, an antibody staining for the enzyme glutamic acid decarboxylase (GAD), which is necessary for the production of GABA, can be used to identify some of the inhibitory synapses in that tissue.

Various other methods have been used to identify possible neurotransmitters, for example: injecting the putative transmitters on to neurons while recording from them, microassays to determine their level in the tissue, labeling of high affinity uptake systems, etc. Each technique has limitations on what it can show. At the moment it is difficult to identify the transmitters involved and their postsynaptic effects at most synapses in the central nervous system. That said, we can make a tentative list of possible generalizations about synapses, although most of them are only supported by our ignorance:

- No axon makes Type I synapses at some sites while making Type II at others.
- No axon in the mammalian brain has been shown to release two different *nonpeptide* neurotransmitters. (But it seems likely that many neurons, including cortical neurons, may release a "conventional" transmitter and a neuropeptide, or in some cases two or more neuropeptides.)
- There is no evidence *so far* in the mammalian brain that the same axon can cause excitation and inhibition at different synapses, but this is certainly possible since the effect of a given transmitter ultimately depends on the kinds of receptors present and their associated ion channels.

Peptides

A remarkable discovery over the last ten years or so has been the existence of many distinct peptides, of various sorts and sizes, which can act as neurotransmitters (see Iverson, 1984, for review). There are, however, reasons to suspect that peptides are different from more conventional transmitters such as acetylcholine or norepinephrine:

- Peptides appear to "modulate" synaptic function rather than to activate it by themselves.
- The action of peptides, in the few cases studied, usually appears to come on slowly and to persist for some time. That is, for times up to seconds or even minutes rather than for a few milliseconds or less as is the case for conventional transmitters.
- In some cases it has been shown that peptides act not at their place of release but at some distance away. This distance may be perhaps some tens of micra or further if carried by a vascular system (as in the path from the hypothalamus to the pituitary). Diffusion takes time. The slow time of onset would be compatible with the possible time delays produced by diffusion.
- There are many examples now known of a single neuron producing (and presumably releasing) more than one neuropeptide.

It has been argued that peptides form a second, slower means of communication between neurons that is more economical than using extra neurons for this purpose. Different peptides are used in the same tissue to enable this communication to have some degree of specificity. (We should remark that so far very little is known about either the *receptors* for peptides or the physiological role of most neuropeptides.)

THE CEREBRAL CORTEX

We shall assume that the reader has some familiarity with the structure of the cerebral cortex and with the behavior of the neurons it contains. For an excellent review of the functional architecture of the primary visual cortex, see Hubel and Wiesel, 1977. This section aims to expand that knowledge. We shall not deal here with the non-neuronal

cells in the cortex (the glial cells, which may outnumber the neurons by 10-50 times) nor with its blood supply, though both these topics are of considerable practical and clinical importance.

The cerebral cortex is conventionally divided into the allocortex (comprising olfactory and limbic cortical areas) and the phylogenetically more recent neocortex, which is all the rest. We shall be concerned almost exclusively with the neocortex, the extensive development of which is characteristic of the mammalian brain, especially the behaviorally more interesting primates.

General Organization

The neocortex consists of two distinct sheets of neurons, one on each side of the head. Each sheet is relatively thin (typical thicknesses run from 1.5 to about 5 mm) and continuous. This is illustrated in Figure 4. Although it is highly convoluted in most of the larger mammals, the neocortex has no slits in it and, as far as we know, no insulating barriers within it. Since it is a continuous finite sheet, it must have an edge. This edge is surrounded by allocortical areas and by various noncortical structures. The sheets on either side of the head are connected by a massive fiber bundle, the corpus callosum. In humans, each sheet has an area of roughly 1000 cm^2 (Henneberg, 1910). In the macaque monkey the figure is nearer 100 cm^2.

Each sheet of the neocortex is highly stratified. An example of the stratification in a typical cortical area is shown in Figure 5A. Historically and didactically, the neocortex has been subdivided into six layers (Lewis, 1878), although a more convenient parcellation can be made into four main layers, which can then be divided further. These four layers are listed below, along with their most prominent features.

- *A superficial layer* (usually referred to as layer I). This layer has rather few cell bodies and consists mainly of axons and apical dendrites. (The presence of this superficial cell-poor layer seems to be characteristic of a "cortical" arrangement of neurons, be it the neocortex, the allocortex, or the cerebellar cortex.)

- *An upper layer* (layers II and III). This layer contains the smaller pyramidal neurons which send their main axons to other cortical areas, either in the same hemisphere or on the opposite side.

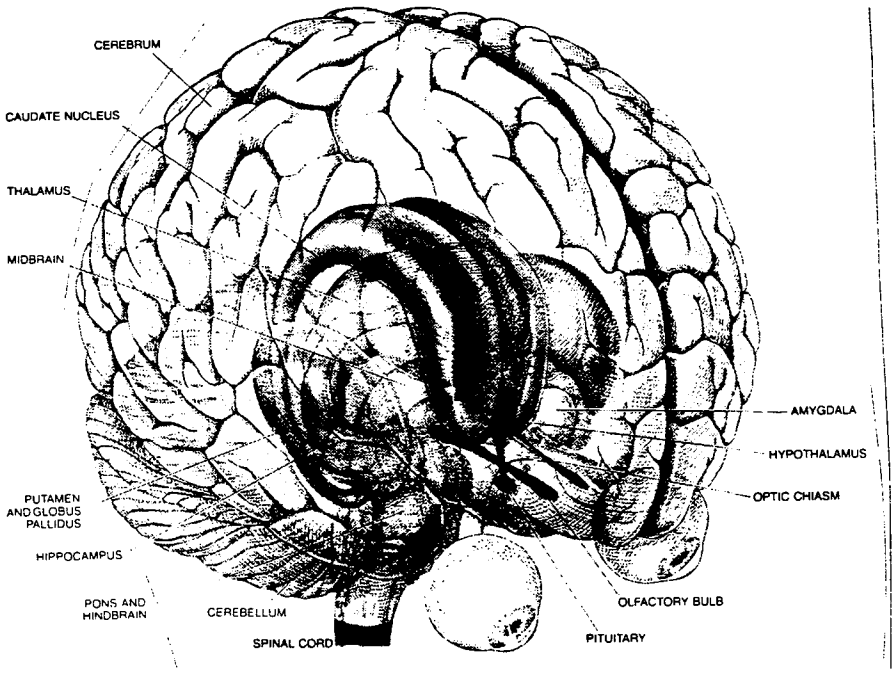

FIGURE 4. A human brain is sliced and opened like a book to demonstrate the continuity of each cortical sheet and its relation to some internal structures. (From "The Organization of the Brain" by W. J. H. Nauta and M. Feirtag, 1979, *Scientific American*, *241*, p. 92. Copyright 1979 by W. H. Freeman & Co. Reprinted by permission.)

- *A middle layer* (layer IV). In this layer are found the densely packed small stellate neurons whose axons commonly ascend vertically to terminate in the upper layers.

- *A deep layer* (layers V and VI). This layer contains the larger pyramidal neurons whose axons leave the cortex to terminate in subcortical structures such as the striatum, the claustrum, the thalamus, the brain stem, and the spinal cord. (Occasional pyramidal neurons are present in this layer which project to other cortical areas rather than projecting subcortically.)

This broad division covers all parts of the neocortex, but there is considerable regional variation in the relative amount of each layer. The middle layer in the primary sensory areas is usually rather thick, e.g., in the striate cortex of primates the middle layer is so pronounced and differentiated that it can be divided into four sublayers (Figure 5B).

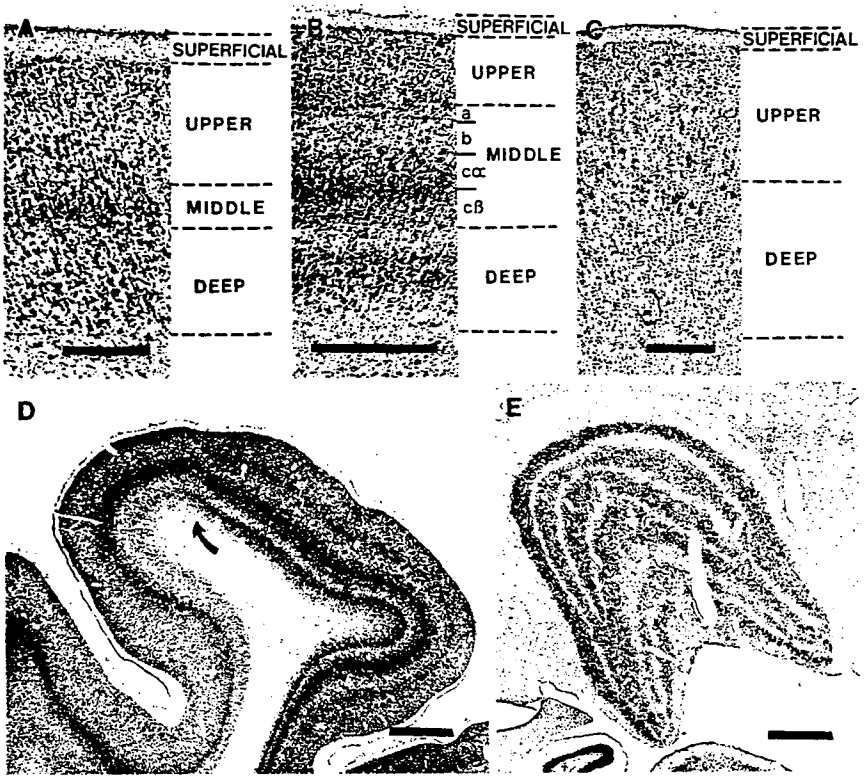

FIGURE 5. *Top*: Some examples of variations in cortical stratification patterns. The stain used is selective for cell bodies. The surface of the brain is at the top in each of these photomicrographs. *A*: Parietal cortex. In most cortical areas, the four main layers are easily recognized. *B*: Striate cortex. A marked differentiation of the middle layer is evident in primary sensory areas. *C*: Motor cortex. The middle layer is virtually absent in the primary motor cortex. *Bottom*: Stains selective for cell bodies are often used to differentiate cortical areas and thalamic nuclei. *D*: Cross-section of the junction between the striate cortex and its immediately adjacent area (area 18). The border is clearly evident (indicated by the arrow) due to the marked differentiation of the middle layer in the striate cortex (right of arrow), and the lack of such a differentiation in the middle layer of area 18 (left of arrow). *E*: The lateral geniculate nucleus is a laminated nucleus, which can easily be identified in cross-sections of the thalamus. Six distinct sheets of neurons can be recognized in the macaque and human lateral geniculate nucleus. All photomicrographs are taken from macaque monkey brains. Bars represent ½ millimeter in A-C, and 1 millimeter in D and E.

In contrast, the middle layer is virtually nonexistent in the primary motor area (Figure 5C).

In addition to the horizontal stratification of neuronal cell bodies, there is a pronounced vertical arrangement of dendritic and axonal

arborizations in the neocortex (Figure 6). Not only do most of the incoming and outgoing axons travel vertically across the layers to enter or exit from the deep aspect of the cortical sheet, but many of the dendritic and axonal processes of neurons in the neocortex are vertically oriented (the ascending dendrites of pyramidal cells are particularly good examples of this—see Figure 14A,B).

The number of neurons per unit *volume* of the neocortex varies somewhat, but the total number of neurons underlying a given unit of surface *area* is remarkably constant from one area of the cortex to another and from species to species (Rockel, Hiorns, & Powell, 1980). In the unshrunken state, this figure is about 80,000 per mm^2 (Powell & Hendrickson, 1981). An exception is the striate cortex of primates, where the figure is about 2½ times as large (Rockel et al., 1980). The reasons for this regularity (and the exception) are not known.

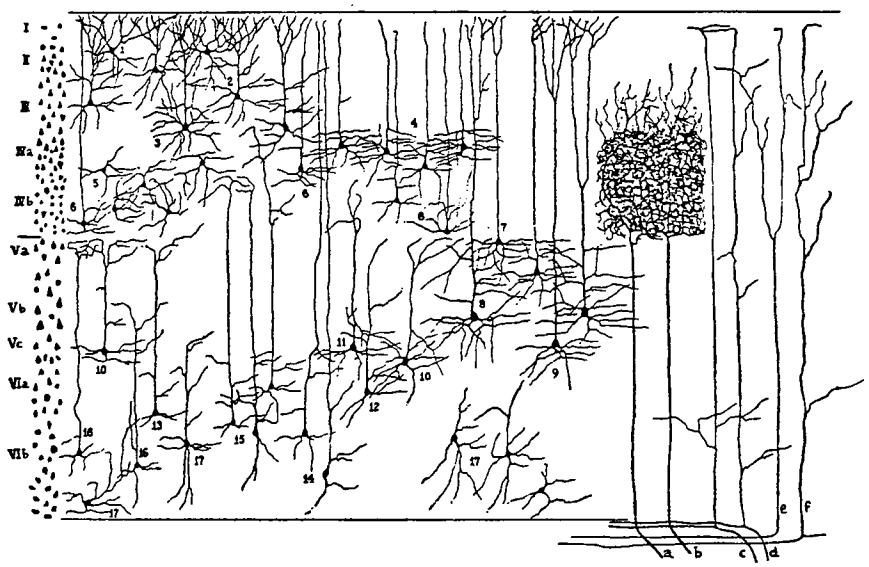

FIGURE 6. The pronounced vertical orientation of many of the dendritic and axonal processes in the neocortex is evident in this diagram of the parietal cortex of an adult mouse. At the left is a diagrammatic representation of all neuronal cell bodies within one very thin section; at the center are the cell bodies and dendrites of some pyramidal neurons, and at the right are some different types of cortical input axons. The surface of the brain is at the top. (From "Cerebral Cortex: Architecture, Intracortical Connections, Motor Projections" by R. Lorente de Nó. In *Physiology of the Nervous System*, p. 282, edited by J. F. Fulton, 1943, New York: Oxford University Press. Copyright 1943 by Oxford University Press. Reprinted by permission.)

Cortical Areas

The neocortex, as already implied, appears to consist of several distinct areas, or "fields" (Rose, 1949). These differ somewhat in their histological appearance[1] (the striate cortex, for example, can be easily recognized in cross section by the presence of a distinct stripe through the middle layer [Figure 5D], although most areas are not so easily recognized.), anatomical connections, and the functions they perform. The hope is that in time it will be possible to parcel out unambiguously the entire neocortex into a number of distinct functional areas. Within each such area we may expect there to be considerable homogeneity of cell types, connections, and functions, all of which are likely to change rather abruptly when one crosses the border of each area and passes into another area. The number of distinct areas in the neocortex of humans (on one side) is likely to be of the order of 100. Presently, the most commonly accepted cortical parcellation scheme is the one that was established by Brodmann (1909) and is illustrated in Figure 7. Although extremely accurate in certain places, this map will undoubtedly be refined in future years.

It has yet to be shown that this simple concept of cortical area may not break down in parts of the neocortex. *If* it holds up, we should be able to count their exact number, so that we could say that in humans there are, say, 137 and not 136 distinct cortical areas. Eventually it should be possible to distinguish each area and thus construct a four-color map of the cortex.

This concept of cortical area appears to hold up fairly well in those cortical areas concerned with early visual processing. In primates there appear to be at least ten of them, covering the region occupied by Brodmann's areas 17, 18, and 19 (Figure 8) It applies very well to the striate cortex (area 17), sometimes called the first visual area (or VI), and to the area known as MT (or the middle temporal area). In the macaque, VI is an exceptionally large area, whereas MT is rather small, being less than 10% the size of VI (Van Essen, Maunsell, & Bixby, 1981; Weller & Kaas, 1983). The size of the other early visual areas will probably fall between these limits. It is important to not lose sight of the basic definition of a cortical area: ridiculously small subdivisions that do not reflect real functional differences can obscure the utility of this concept.

---

[1] The common terms used for these differences are cytoarchitectonics and myeloarchitectonics. The former refers to the differences in neuronal density and to the relative development of individual cortical layers; the latter refers to differences in the distributions of axons (especially myelinated axons) within the cortex which varies from area to area.

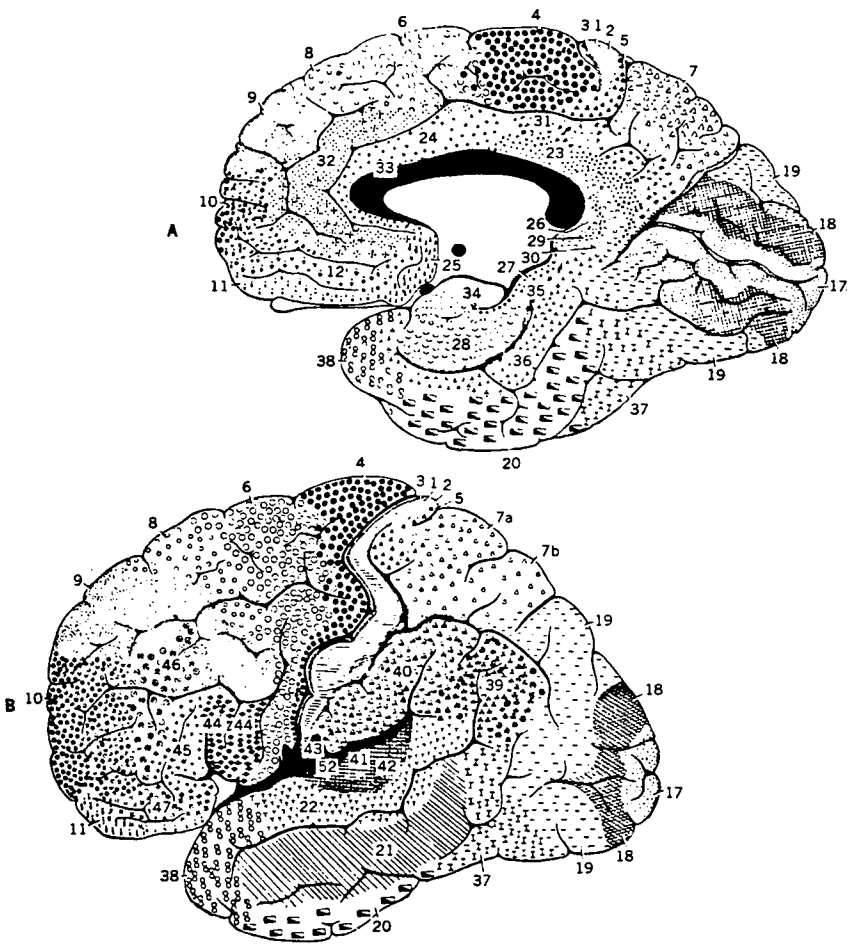

FIGURE 7. Brodmann's areas of the human cerebral cortex. Each of his areas are numbered and indicated by different symbols. *A*: Medial surface of the cerebral cortex (the black areas are occupied by fiber bundles crossing the midline to connect the two hemispheres). *B*: Lateral surface of the cerebral cortex. (From *Vergleichende Localisationslehre der Grosshirnrinde in Ihren Prinzipien Dargestellt auf Grund des Zellenbaues [Principles of comparative localization in the cerebral cortex presented on the basis of cytoarchitecture]*, by K. Brodmann, 1909, Leipzig: Barth. Copyright 1909 by Barth Publishing. Reprinted by permission.)

Cortical Inputs

An important feature of the neocortex is that almost all the outside information it receives (either from the sensory periphery or from other subcortical centers), with the exception of some olfactory

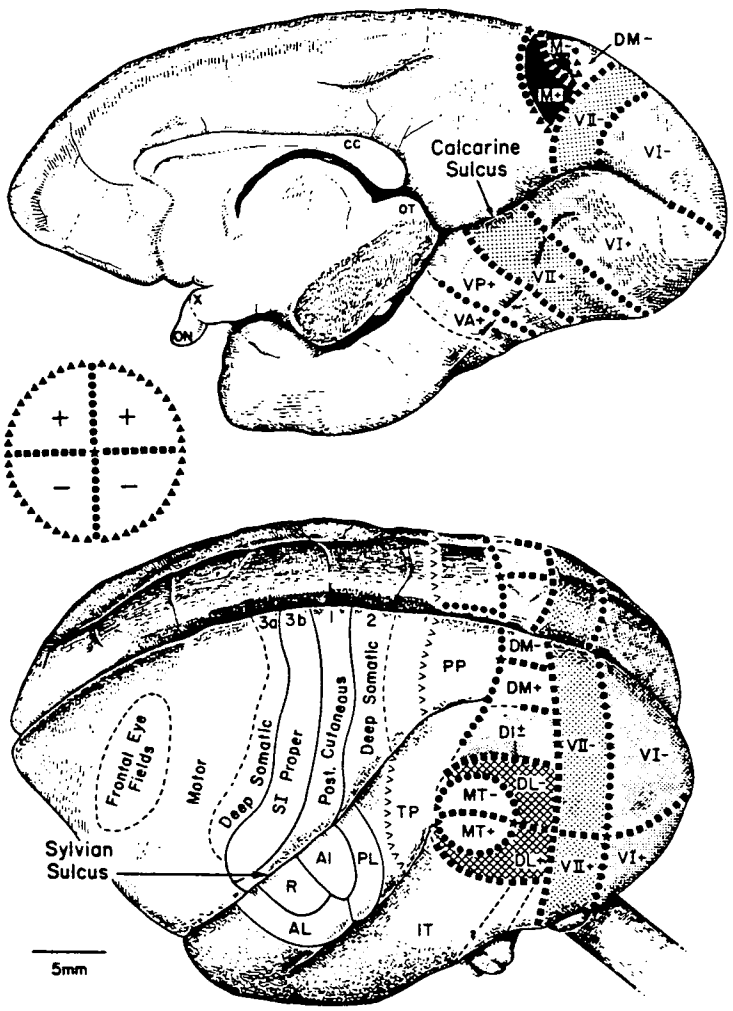

FIGURE 8. Visual processing areas of the owl monkey cerebral cortex. Each of the different types of shading represent different areas. +'s indicate the regions representing the dorsal half of the visual field. −'s indicate the ventral half of the visual field. DI, dorsointermediate visual area; DL, dorsolateral crescent visual area; DM, dorsomedial visual area; IT, inferotemporal cortex; M, medial visual area; MT, middle temporal visual area; PP, posterior parietal cortex; VA, ventral anterior visual area; VP, ventral posterior visual area; V1, first visual area; V2, second visual area. (From "Visual Response Properties of Neurons in Four Extrastriate Visual Areas of the Owl Monkey (*Aotus trivirgatus*): A Quantitative Comparison of the Medial, Dorsomedial, Dorsolateral, and Middle Temporal Areas" by J. F. Baker, S. E. Petersen, W. T. Newsome, and J. Allman, 1981, *Journal of Neurophysiology*, *45*, p. 400. Copyright 1981 by The American Physiological Society. Reprinted by permission.)

information, passes through the thalamus. Input systems terminate on neurons in the thalamus, and these thalamic neurons, in turn, project to the cerebral cortex (Figures 9A and 10). Though the terminations of thalamic axons may account for only a small proportion of the total synapses in any given cortical area,[2] the thalamus is clearly the major

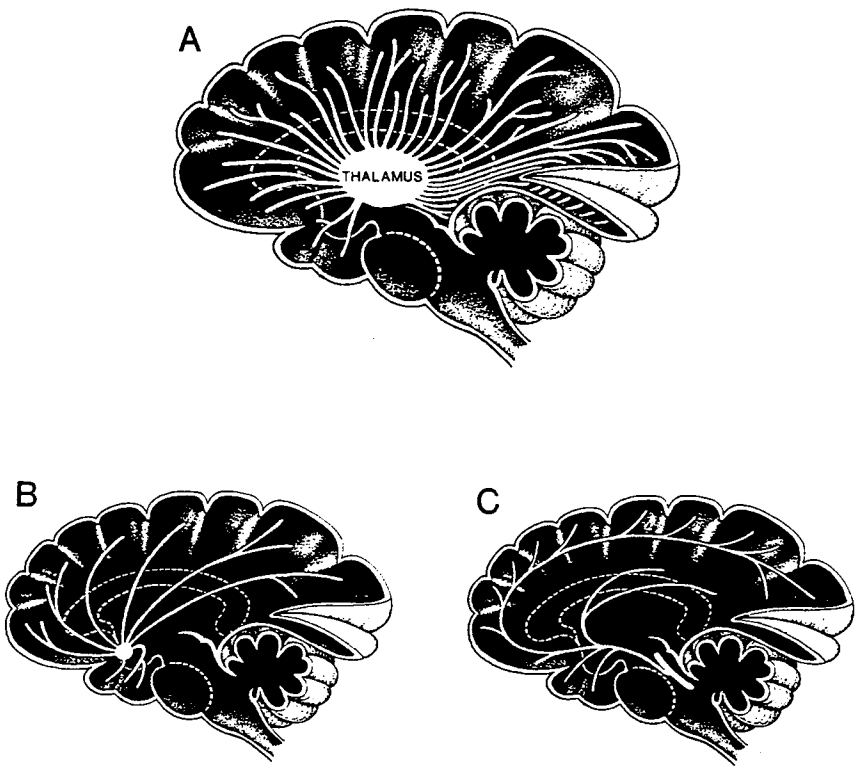

FIGURE 9. Some of the inputs to the neocortex. *A*: Most of the information entering the neocortex gets there through the thalamus. *B*: A diffuse cholinergic input arises in the basal forebrain. *C*: Diffuse noradrenergic and serotonergic inputs arise in the brain stem.

[2] Recent synapse counts indicate that in the monkey striate cortex, approximately 35% of the total synaptic population comprises middle layer synapses (O'Kusky & Colonnier, 1982). Reported percentages of thalamocortical synapses within the middle layer of the striate cortex range from 5% (Garey & Powell, 1971) to 29% (Tigges & Tigges, 1979). These data suggest that thalamocortical synapses account for 2-10% of the total synaptic population in the striate cortex, but this calculation does not take into account the thalamocortical synapses which terminate outside the middle layer (e.g., in layer I and in layer VI).

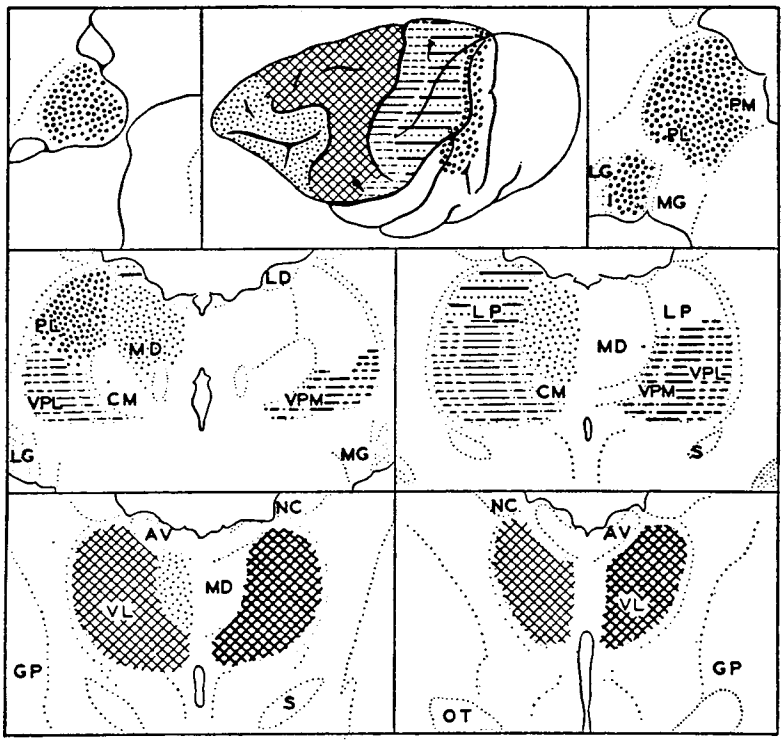

FIGURE 10. A schematic diagram demonstrating part of the systematic relationship of the thalamus with the cerebral cortex and the thalamic termination of several input systems in macaque monkeys. The cerebral cortex (top center) is viewed from the left side, and several frontal cross-sections of the thalamus are illustrated, with the caudalmost section at top left and the rostralmost section at bottom right. Similarly marked parts on the left sides of the thalamic sections and the cerebral cortex have thalamocortical connections. The terminations of afferent systems are represented on the right sides of the thalamic sections as follows: inputs from the cerebellum by heavy cross-hatching, somatic sensory inputs from the leg area by short heavy dashes, somatic sensory inputs from the arm area by heavy lines, and somatic sensory inputs from the face by heavy dots and dashes. Abbreviations in the thalamic sections indicate histologically identifiable nuclei. (From *The Primate Thalamus*, p. 189, by A. E. Walker, 1938, Chicago: University of Chicago Press. Copyright 1938 by the University of Chicago Press. Reprinted by permission.)

center through which the cerebral cortex has access to outside information. The thalamus is, therefore, often referred to as the "gateway" to the cerebral cortex.

There are a few exceptions to this general rule. The following pathways do not relay through the thalamus:

- Diffuse innervations to the cortex arise in a number of brain stem and basal forebrain areas (Figures 9B,C). Among them are a noradrenergic system that arises principally in the locus coeruleus (Andén et al., 1966), a serotonergic system that arises mainly in the dorsal raphé nucleus (Bobillier et al., 1976), and a cholinergic system that arises in the basal forebrain (Mesulam & Van Hoesen, 1976). These systems spread throughout the cortex both horizontally and vertically and do not appear to be organized with any topographic finesse. In addition to these, it has been suggested recently that there may be a diffuse GABAergic system innervating the neocortex, which arises in the hypothalamus (Vincent, Hökfelt, Skirboll, & Wu, 1983). The exact function of these diffuse inputs is not known, but it is important for theorists to be aware of their existence.

- A structure called the claustrum—situated deep to the insular region of the cortex, receives inputs from (Carman, Cowan, & Powell, 1964) and projects to (Macchi, Bentivoglio, Minciacchi, & Molinari, 1981) almost all areas of the cortex. Since, apart from its diffuse innervation from the brain stem, it receives no other input, it could well be described as a satellite of the cortex. Only the visually responsive part of the claustrum has been intensively studied, and it has been shown to be systematically connected with the striate cortex and the adjacent visual area (LeVay & Sherk, 1981a).

- Restricted neocortical projections arise in the hippocampus (Swanson, 1981) and amygdala (Porrino, Crane, & Goldman-Rakic, 1981). These tend to terminate most heavily in cortical areas that are removed from the primary sensory and primary motor areas.

The thalamus comprises a number of *specific nuclei*; these have well-defined inputs and project to restricted portions of the cerebral cortex. The thalamic nuclei, like the cortical fields, can be differentiated in terms of histological appearance, connections, and function. The lateral geniculate nucleus, for example, is the relay center for inputs from the retina that pass on to the striate cortex. It is a distinct, laminated nucleus which can be easily differentiated histologically from the surrounding structures (Figure 5E). Neurons in the lateral geniculate nucleus receive visual signals from the axons of ganglion cells in the retina and project, in turn, to the striate cortex.

The inputs to the cortex from the thalamus terminate primarily in the middle layer (Cajal, 1911). Where the middle layer is sparse or

absent (as in the motor cortex), they usually terminate in the lower part of the upper layer (Jones, 1975a). In the electron microscope, the terminal boutons of thalamocortical axons are all Type I (Jones & Powell, 1970a) and are known to be excitatory. The axons of the thalamic neurons that project to the cortex rarely have collaterals within the main body of the thalamus (Friedlander, Lin, Stanford, & Sherman, 1981).

In addition to the specific nuclei, there are a number of less specific (sometimes called nonspecific) thalamic nuclei. The most prominent group of the less-specific nuclei is known as the intralaminar nuclei and occupies a thin, vertical, sheet-like zone extending anteroposteriorly through the center of the thalamus. While most of these are rather small, one—the centromédian nucleus—is very prominent in the human brain. The neurons of the intralaminar nuclei project both to the striatum and to the cerebral cortex (Jones & Leavitt, 1974; Powell & Cowan, 1967). Their cortical projection, instead of terminating in the middle layer of the cortex, terminates mainly in the superficial layer (layer I). Moreover, the cortical projections of the intralaminar nuclei are not confined to a single cortical field, but tend to be rather widespread. Our present knowledge of the less-specific thalamic nuclei and their role in cortical function is quite vague.

Each of the specific thalamic projections to the cortex is accompanied by a reverse projection from the cortex to the thalamus. The spatial organization of the reverse projection reciprocates, fairly precisely, the forward one. This reverse projection arises from cells at the bottom of the deep layer and terminates directly on the peripheral dendrites of thalamocortical relay cells. Their terminations are also Type I (Guillery, 1969; Szentágothai, Hámori, & Tömböl, 1966). Although they are very numerous, the function of these reverse projections is not known.

A very remarkable nucleus, called the thalamic reticular nucleus forms a thin shell around the main body of the thalamus. It is only a few cells thick. Its neurons are very large, with protrusions on their dendrites called spines (M. E. Scheibel & A. B. Scheibel, 1966). This nucleus does *not* project to the cortex but projects back into the thalamus (Jones, 1975b). It appears to receive small collaterals from most of the axons that pass between the thalamus and the cortex. It also gets some input from the traffic between the thalamus and the striatum. Its axons are inhibitory (Houser, Vaughn, Barber, & Roberts, 1980; Montero & Scott, 1981), and they have extensive axon collaterals *within* the reticular nucleus (Scheibel & Scheibel, 1966). Obviously this nucleus occupies a very strategic place in the brain. It deserves more attention, both from experimentalists and theorists (Crick, 1984).

In addition to their thalamic inputs, most cortical neurons receive inputs from other cortical areas either in the same hemisphere (in

which case they are called associational inputs) or in the opposite hemisphere (where they are known as commissural inputs). It is important to note that a typical cortical area is not connected directly to all or even most other cortical areas. Usually it projects to a handful of other areas, although the areas that are removed from the primary sensory or motor areas tend to project more widely (Jones & Powell, 1970b; Pandya & Kuypers, 1969). But if a cortical area projects to another cortical area, its projections are usually topographically organized, at least on a coarse scale. That is, as far as position in the sheet is concerned, connections between areas are not random; neighboring neurons in a field tend to project to neighboring regions in other fields in some systematic way. Moreover, as a general rule, projections from one field are usually matched by a reciprocal projection from that field which is also topographically organized. To a considerable extent the forward and backward mappings coincide, at least on a coarse scale, but are not symmetrical in all details. Rockland and Pandya (1979) and Maunsell and Van Essen (1983) suggest that for the early visual areas, a forward projection (forward with respect to the retina) is likely to project predominantly into the middle layer, whereas the reverse projection is likely to avoid the middle layer and instead terminate largely in the superficial and deep layers (Figure 11).

Topographically organized maps occur in many cortical areas. Anatomical and electrophysiological studies show that in most cortical areas there is a more or less topographic representation of the periphery upon the cortical surface. Detailed maps of cortical representation patterns are available for the areas concerned with sensory input (and motor output). As might be expected, such maps vary somewhat from individual to individual. Their broad topography is usually not linear; the extent of cortex representing a given region of the periphery is roughly proportional to the peripheral innervation density. Thus, the fovea is heavily over-represented in the striate cortex (Figure 12A), and the hand occupies a bigger region than the trunk in the somatic sensory (Nelson, Sur, Felleman, & Kass, 1980; C. N. Woolsey, Marshall, & Bard, 1942) and motor cortical areas (C. N. Woolsey et al., 1952; see Figure 12B).

All this makes good sense. Because neurons in the cortex cannot act directly on more distant neurons in the same area (see later), such "maps" bring into proximity neurons which are likely to have related inputs. They also allow more space for the more densely innervated regions.

This systematic mapping of the periphery is likely to be a feature of many cortical areas, not merely those near the sensory input. In general terms, as one proceeds further from the sensory input, the

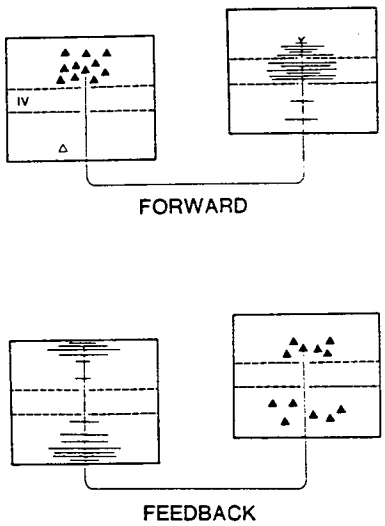

FIGURE 11. A schematic diagram of the characteristic laminar distributions of cell bodies and terminals in forward and feedback cortico-cortical pathways. Forward pathways arise mainly from the upper layer and terminate mainly in the middle layer. Feedback pathways arise from both upper and lower layers and terminate mainly outside the middle layer. Triangles represent cell bodies, and axons are represented by thin lines. (From "The Connections of the Middle Temporal Visual Area and Their Relation to a Cortical Hierarchy in the Macaque Monkey" by J. H. R. Maunsell and D. C. Van Essen, 1983, *Journal of Neuroscience*, *3*, p. 2579. Copyright 1983 by The Society for Neuroscience. Reprinted by permission.)

mapping of the periphery becomes more diffuse. At the same time the neurons respond to more elaborate "features" of the input.

There are irregularities within a cortical area. It should not be inferred from the above description that cortical areas are relatively uniform. On the contrary, on the scale of 1 mm to 2 mm, patchiness is almost always the rule. The patchiness can, in many instances, be attributed to disjunctive distributions of inputs. Dense foci of terminations, separated by zones that are relatively free of terminations, have been described for the inputs from the thalamus and from other cortical areas.

This is especially clear in cases where the inputs are not continuous but discrete. For example, the thalamic input from either of the two eyes form nonoverlapping stripe patterns (not unlike fingerprints when viewed from the surface) in the middle layer of the striate cortex (Figure 13) (Hubel, Wiesel, & LeVay, 1977; LeVay, Connolly, Houde, & Van Essen, 1985). In the somatic sensory system of rodents, inputs

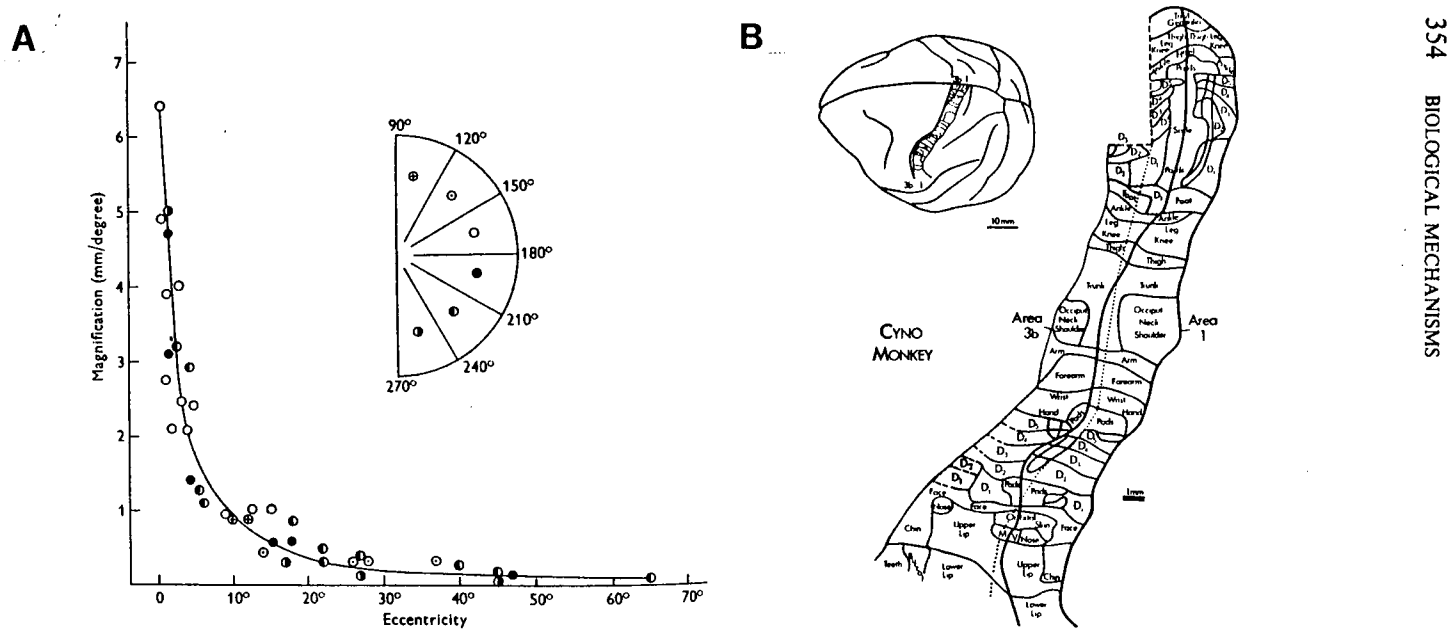

FIGURE 12. Though continuous, the cortical representations of the periphery are often not faithful duplications of the periphery. *A*: In the primary visual cortex, the fovea occupies a disproportionately large percentage of the total area. The amount of cortical surface area devoted to a degree of visual field is plotted as a function of retinal eccentricity. (From "The Representation of the Visual Field on the Cerebral Cortex in Monkeys" by P. M. Daniel and D. Whitteridge, 1961, *Journal of Physiology*, *159*, p. 212. Copyright 1961 by Cambridge University Press. Reprinted by permission.) *B*: In the primary somatic sensory cortex of primates, the representation of the hands (in particular, the digits) occupies a bigger region than the trunk. The areas of the cortex devoted to each body part are depicted on a flattened map of the primary somatic sensory cortex. (From "Representations of the Body Surface in Postcentral Parietal Cortex of *Macaca fascicularis*" by R. J. Nelson, M. Sur, D. J. Felleman, and J. H. Kaas, 1981, *Journal of Comparative Neurology*, *192*, p. 614. Copyright 1981 by Alan R. Liss, Inc. Reprinted by permission.) The data shown in both A and B are derived from macaque monkeys.

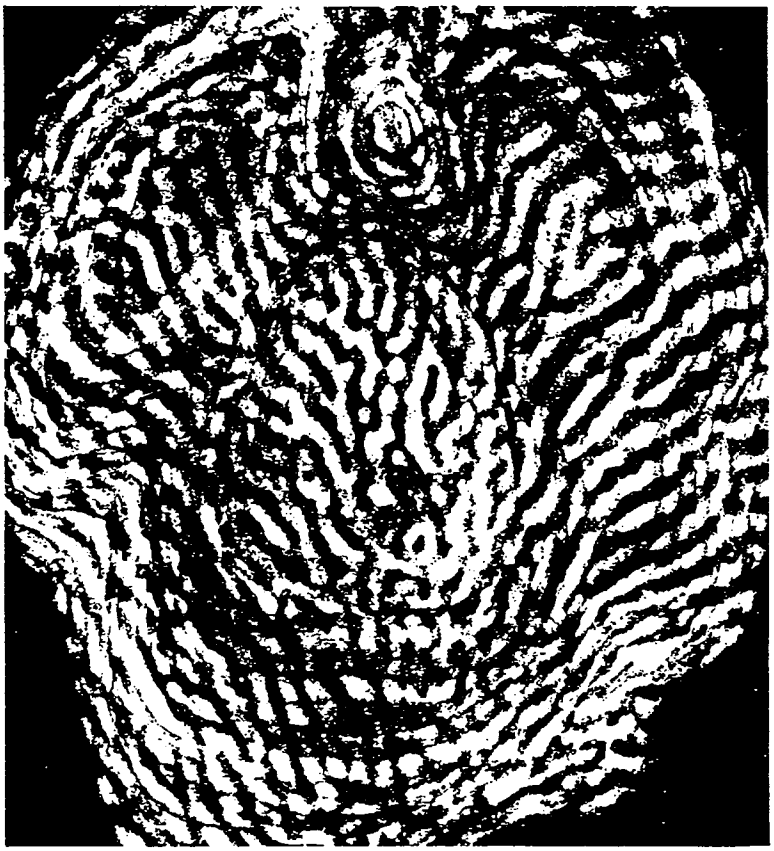

FIGURE 13. A surface view of the middle layer of the striate cortex in a macaque monkey in which the input from one of the eyes is labeled autoradiographically with tritium. In this figure, the label, in the form of exposed photographic silver grains, appears light against a dark background. The disjunctive, stripe-like patterns of input from each eye is clearly evident. Due to the curvature of the cortex, this figure is a montage of many photomicrographs. (From "The Complete Pattern of Ocular Dominance Stripes in the Striate Cortex and Visual Field of the Macaque Monkey" by S. LeVay, M. Connolly, J. Houde, and D. C. Van Essen, 1985, *Journal of Neuroscience*, *5*. Copyright 1985 by The Society for Neuroscience. Reprinted by permission.)

from each mystacial vibrissa project, through independent channels in the thalamus, to separate foci in the middle layer of the somatic sensory cortex called "barrels" (T. A. Woolsey, 1978). Neural connections *between* barrels appear less strong than those *within* each barrel. Similar focal arrangements have been proposed for the thalamic input to other cortical areas.

Disjunctive, stripe-like patterns have also been demonstrated in certain terminations coming from other cortical areas. For example, the projection from the striate cortex to other visual areas is usually

irregular. Stripe-like patterns of terminations have been described for their projections to area 18 (Gilbert & Wiesel, 1981; Maunsell, Newsome & Van Essen, 1980). In the primary auditory cortex, there are irregular rows of commissural terminations which in cross-section appear as multiple column-like clusters. Such clusters contain neurons that display certain kinds of binaural interaction (Imig & Brugge, 1978).

The various kinds of patches outlined above appear to underlie a tendency of neurons aligned vertically to display certain similarities in functional properties. Other response properties, not necessarily correlating with input signals, are also aligned vertically. For example, in the visual system of monkeys, the cells that project to the cortex have center-surround receptive fields, and the neurons in the upper and lower layers of the striate cortex respond specifically to lines oriented in a particular orientation (Hubel & Wiesel, 1968). Furthermore, neurons with similar orientation preferences are often aligned vertically, though this may not always be true. However, in a general sense it has been found that when recordings are made with a microelectrode inserted perpendicular to the surface of the cortex, the response properties of all the neurons encountered throughout the depths of the cortex tend to be similar.

These various types of patchinesses have suggested to some authors that the cortex is organized on a "modular" basis, either involving columns about ½ mm or so in diameter or into minicolumns, some 50 microns or so across (Mountcastle, 1978). In the cortex there is really no evidence for true modules of the type found, for example, in the fly's eye and rather little for more irregular modules. The exact basis on which the cortex is organized remains to be discovered.

Cortical Outputs

All the axons that leave a cortical area make Type I (excitatory) synapse in their projection fields. The subcortical projections of the cortex, in contrast to the subcortical inputs, are not subject to the constraint of always relaying through the thalamus. Thus, the cortex can project directly to the spinal cord or superior colliculus, but the spinal and collicular inputs to the cortex always relay in the thalamus.

As mentioned earlier, there is a close relationship between populations of cortical output neurons and their laminar positions. This is particularly evident for the populations of neurons that project subcortically. For example, the cortical neurons which project to the thalamus are situated at the bottom of the deep layer, whereas those that project to the striatum and brain stem tend to be concentrated in more

superficial portions of the deep layer (Jones & Wise, 1977; J. S. Lund, R. D. Lund, Hendrickson, Bunt & Fuchs, 1975).

As with the input systems, most of the output systems of the neocortex are patchy. Patches of neurons have been demonstrated to project to other cortical areas and others to project subcortically (Jones & Wise, 1977; Murray & Coulter, 1981). Although the groups of cells giving rise to the various projections are segregated into layers, physiological studies indicate that neurons producing *similar* outputs are vertically aligned. This has been most convincingly demonstrated in the motor cortex, where microstimulation points producing isolated contractions of given muscles have been shown to be organized in cylindrical zones which extend throughout the depths of the cortex (Asanuma & Rosén, 1972).

THE NATURE OF NEOCORTICAL NEURONS

Experimental Methods

The shape of neurons in many areas of the brain has traditionally been studied in Golgi-stained material. The Golgi method stains, in their entireties, a few cells here and there, apparently "at random." The haphazard selectivity of this method allows investigators to visualize clearly most of the dendritic processes of individual neurons. Unfortunately, though, the technique is capricious and does not always stain the axon and axon collaterals completely. A more recent method, using intracellular injections of an appropriate marker such as the fluorescent dyes procion yellow or lucifer yellow or the enzyme horseradish peroxidase (HRP), is capable of producing a similar picture with better reliability and control, but is not suited for sampling large populations of neurons since each neuron must be separately injected.

Very fine details, such as the morphology of the synapse, cannot be seen using the light microscope. The electron microscope is necessary for this. The extremely high magnifications needed to differentiate such details, however, make it difficult to reconstruct (using a very large number of serial sections) even a very small neuron in its entirety. (This problem can be partially alleviated by combining standard electron microscopic procedures with the Golgi method or a variety of labeling methods.)

It is important to realize that each of the techniques presently available is limited in its sampling capacities. Many of the following conclusions, therefore, are rather general and apply to broad groupings of

cortical neurons. Except in rare cases, there is little quantitative information on specific cell types.

Cell Types in the Neocortex

The most common and most characteristic neurons of the cerebral cortex are the pyramidal cells (Figure 14A). Such neurons, as their name implies, have a pyramid-shaped cell body. A thick apical dendrite ascends vertically from the apex of the cell body. It often has tufts of branches that arborize in specific laminae. A number of basal dendrites also extend out from the bottom of the cell; these branches spread horizontally for some distance. All dendrites of pyramidal cells possess large numbers of dendritic spines. The axons of pyramidal cells emerge either from the bottom of the cell body or from one of the basal dendrites and generally leave the area.

Though these basic features are common to all pyramidal cells, there are a variety of different types of pyramidal cells (Figure 14B). On the basis of cell body size alone, they range from about 10 microns up to 100 microns in diameter (Cajal, 1911; Conel, 1939). The branching patterns of the apical and the basal dendrites also differ tremendously, as do the target sites of the projection axons and the arborizations of axon collaterals which participate in local circuits. These variations do not appear to be random, and a number of classes of pyramidal cells have been described (Cajal, 1911; Lorente de Nó, 1943; J. S. Lund, 1973). Unfortunately, the information presently available for differentiating the various types of pyramidal cells is far from complete. It is unlikely that there are as few as ten distinct types: There could be as many as a thousand.

All axons that leave the cortex belong to pyramidal cells. The converse, however, may not be true. Though all pyramidal cells were once considered to be projection neurons (Cajal, 1911), recent data suggest otherwise. Using intracellular HRP injections, Gilbert and Wiesel (1983) have found a few pyramidal cells with extensive local axonal arborizations that do not project out of the striate cortex.

The remainder of the neurons in the cerebral cortex can be broadly categorized as "nonpyramidal cells" (Figure 14C). By definition, they comprise all neurons whose cell bodies are not pyramid-shaped. They constitute a large number of cell types which can be differentiated on the basis of the shape of the cell body, size, dendritic arborizations, and axonal ramifications. For example, some have relatively spherical cell bodies with dendrites that radiate in all directions and therefore are star-shaped, or "stellate." Others are spindle-shaped and have dendritic

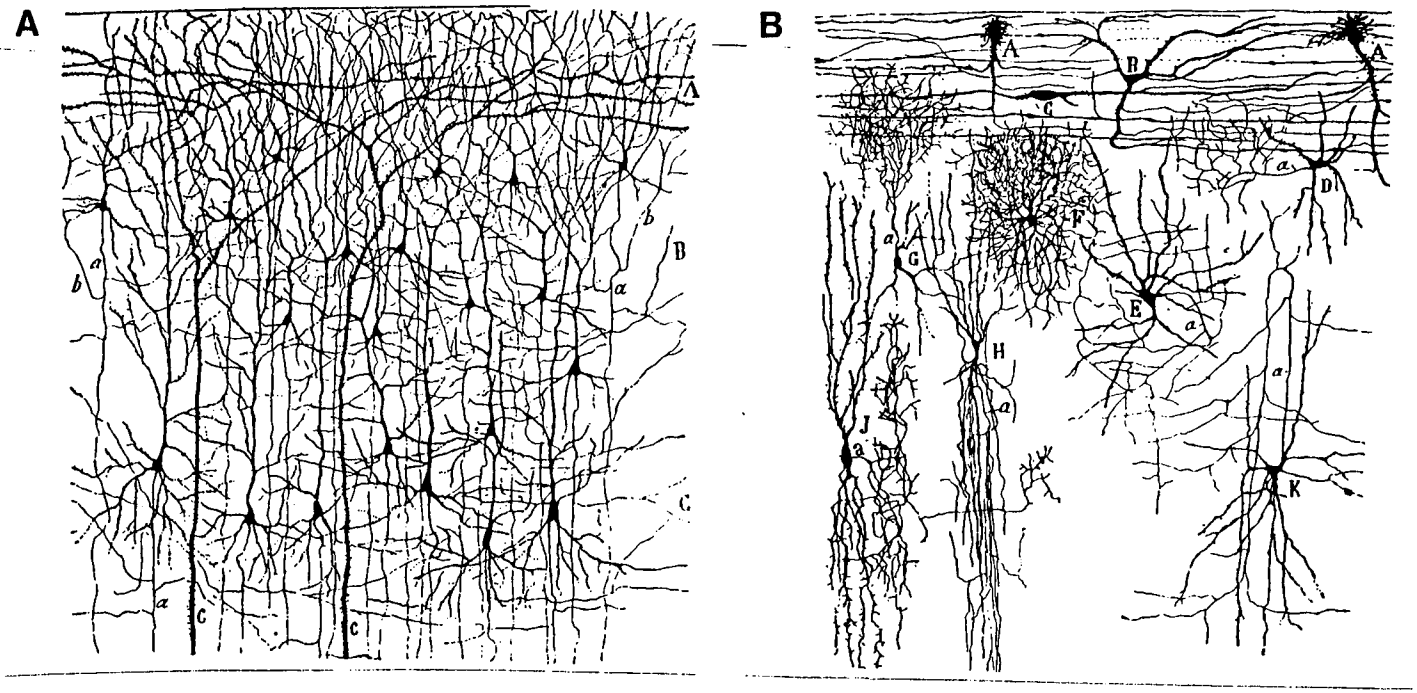

FIGURE 14. Drawings of Golgi-stained neurons of the cerebral cortex. *A*: The most characteristic neurons of the cerebral cortex are the pyramidal cells. They are characterized by a pyramid shaped cell body, a thick apical dendrite, and a number of basal dendrites. *B*: There are a variety of different types of pyramidal cells. *C*: The remainder of neurons in the cerebral cortex can be broadly referred to as nonpyramidal cells. There are a variety of different types of nonpyramidal cells. See text for further details. (From *Histologie du Systeme Nerveux de L'homme et des Vertèbrès, T. 2. [Histology of the Nervous System of Man and of Vertebrates, Vol. 2]*, by S. Ramón y Cajal, 1911, Paris: Maloine. Copyright 1911 by Librairie Maloine. Reprinted by permission.)

arborizations which are aligned vertically. Again, the number of distinct types is not known. It is unlikely to be as few as ten; it could run to several hundreds.

Instead of discussing in detail the varieties of nonpyramidal cells that are known, we should like to draw the reader's attention to two characteristics of nonpyramidal cells that are of significance in the context of this chapter. One is that the nonpyramidal cells are local circuit neurons, viz., they do not project out of a given cortical area. The other is that there are two basic types of nonpyramidal cells according to their dendritic morphology. One type has numerous spines on its dendrites, and the other has few or no spines.

A convenient categorization combines the entire population of pyramidal cells and also the spiny nonpyramidal cells into a group called "neurons with spines," and the remaining nonpyramidal cells into a group called "neurons without many spines." The following generalizations are applicable when cortical neurons are categorized in this way (see Table 1).

Neurons With Spines

The neurons with spines are the pyramidal cells, the star pyramids, and the spiny stellate cells. Such cells receive Type I synapses mainly on their spines. Usually, each spine has only a single Type I synapse;

TABLE 1

	Neurons With Spines	Neurons Without Spines
Input to Spines	Usually one Type I Occasionally one Type I plus one Type II	–
Input to Dendrites	Both types	Both types
Input to Soma	Type II only	Both types
Input to Axon Hillock	Multiple Type II on many pyramidals in layers II and III (see text)	–
Output	Type I	Usually Type II

Working Assumptions: Type I = excitation Type II = inhibition

a minority of spines has an additional Type II synapse. No spine seems to have only a Type II synapse by itself (Colonnier, 1968).

The Type II synapses on spiny cells tend to be fewer in number than Type I synapses. They do, however, occupy sites that are suited for effectively influencing impulse generation by the neuron (which generally takes place at the axon initial segment). They usually occur on the proximal shafts of the dendrites or on the cell bodies (Gray, 1959). A special type of Type II synapse is found on the initial segment of the axon of many pyramidal cells (Peters et al., 1968).

The axon of a spiny cell forms only Type I synapses (LeVay, 1973). Pyramidal cells almost always send the main branch of their axon out into the white matter below the cortex. Spiny stellates almost always have a local axon which does not leave the cortex but ramifies instead in close proximity to its cell of origin.

Neurons Without Many Spines

Although neurons without spines may have transient spines on their dendrites early in development, in the mature state they have relatively few, if any, spines.

Such cells, of which there are several obviously different morphological types, receive both Type I and Type II synapses on their dendritic shafts and on their cell bodies (Gray, 1959). The axons of nonspiny neurons do not descend to the white matter but ramify locally. Their axons are believed, in most cases, to form only Type II synapses (LeVay, 1973), but there may well be exceptions.

It remains to say that spiny cells are in the majority (80%). Of these, a fair number, perhaps as many 25% can be nonpyramidal in some cortical areas. The nonspiny cells are in a minority (20%). Unfortunately these percentages are only very approximate.[3]

---

[3] These approximations are derived from the results of Tömböl (1974), who found 59% of the neurons in the visual cortex of monkeys to be pyramidal neurons, 7.4% to be large stellate neurons, and 33.6% to be small stellate neurons. Of the population of small stellate neurons, 63% were found to occur in the middle layer. These results did not take into account the spiny stellate neurons which also need to be included in the spiny category. Since spiny stellate neurons are small (Cajal, 1911), restricted to the middle layer (J. S. Lund, 1973), and constitute the major component of the middle layer (J. S. Lund, 1973), we have added 21% to the overall percentage of pyramidal neurons reported by Tömböl, to obtain a very rough figure of 80% for our estimate of spiny neurons.

Excitation Versus Inhibition

We shall now discuss these generalizations (summarized in Table 1) on the assumption that most Type I synapses excite and most Type II synapses inhibit. The main point to grasp is the asymmetry between excitation and inhibition.

Excitatory neurons and excitatory synapses are clearly in the majority. Only excitation is sent out of any area of the cortex to other places. On the other hand, the inhibitory synapses are more strategically placed on the cells, being nearer the impulse initiating site at the axon hillock, at least for spiny cells. This is almost a necessity if the system is not to get out of hand, since any area of the cortex feeds back excitation to itself in large amounts. The synapses formed by axons entering the cortex are in a minority, sometimes a very small minority. The great majority come from the axons of other cortical cells: mostly, but not entirely, from those within a few hundred microns.

Inhibition seems to have priority in two ways. Inhibitory neurons have excitatory inputs on their cell bodies, so that they can be brought into action quickly. Excitatory neurons, on the other hand, receive inhibition at strategic sites. This preference for inhibition must be set against the fact that (excluding some of the diffuse innervation) the inputs entering the cortex are all excitatory. Thus any inhibition needed must be generated from this excitation. This requires an extra step and therefore will take time. It seems as if the cortex is arranged so that this time delay is minimized by the regularities discussed above.

In special conditions, such as epilepsy and under the influence of hallucinogenic drugs, the cortical system may go into full-scale oscillation, presumably because the balance between excitation and inhibition has been upset.

Special Cell Types

Nonspiny neurons are of various types. Many of them are stellate. Here we describe three unusual types of nonspiny neurons whose properties may be of special interest.

Chandelier cells. So-called because their axons end in a set of vertically oriented beaded terminal segments which make them look somewhat like chandeliers (Szentágothai & Arbib, 1974). An alternative name for them is "axo-axonic" cells (Somogyi, 1977).

These beaded terminal segments turn out to be the sites of multiple Type II synapses upon the axon initial segments of nearby pyramidal

cells (Somogyi, 1977). Each chandelier cell axon has at least 50 to 200 such terminal segments (Somogyi, Freund, & Cowey, 1982), each of which may form 5 to 10 synapses (Fairén & Valverde, 1980; Somogyi et al., 1982). (There does not appear to be any corresponding cell which makes Type I synapses on axon hillocks.) Each pyramidal cell may receive up to six such terminal segments, presumably from different chandelier cells (Fairén & Valverde, 1980; Somogyi et al., 1982).

It is difficult to resist the impression that a chandelier cell can veto the output of a whole set of neighboring pyramidal cells. Whether this is really true is not known, but the fact that chandelier cells are probably GABAergic (Peters, Proskauer, & Ribak, 1982) and form Type II synapses strongly implies that they are inhibitory.

The number of chandelier cells is not known. A not unreasonable guess would be that they form about 1% of the total neuronal population.[4] It is not known how many types of chandelier cells exist. So far no stain has been found which selectively stains them.

Basket cells. These cells have vertically oriented stem axons which give rise to several horizontally disposed collaterals. The collaterals subsequently give off obliquely or vertically directed fine terminal branches, which are commonly clustered and resemble baskets (Cajal, 1911; Marin-Padilla, 1969), though not all terminal branches resemble baskets (Jones, 1975c). These terminal branches often form loops of boutons around the cell bodies of pyramidal cells (10 to 20 terminal branches may converge upon certain pyramidal cells; Martin, Somogyi, & Whitteridge, 1983). The characteristic fine axonal sprays of these cells are rather difficult to visualize, so it is not certain that basket cells occur in all cortical areas. The terminal boutons of basket cells are Type II in morphology and are thought to contain GABA (Martin et al., 1983; Somogyi, Kisvárday, Martin, & Whitteridge, 1983).

Again, one cannot resist the impression that the basket cell is likely to exert a veto (or at least a partial veto) upon the output of the cortical projection neurons.

Bipolar cells. These are cells whose dendrites and axons form a very long thin bundle in the vertical direction (M. L. Feldman & Peters, 1978). Some of them have been shown to contain peptides such as somatostatin, cholecystokinin (CCK), and vasoactive polypeptide (VIP)

---

[4] This guess is based on three arbitrary assumptions: (a) all pyramidal neurons receive initial segment synapses from chandelier cells, (b) every chandelier cell innervates about 200 pyramidal neurons, and (c) every pyramidal neuron is innervated by about five different chandelier cells. Together, these suggest that there are about 40 pyramidal neurons for each chandelier cell. Since pyramidal neurons account for about 60% of the neurons in the visual cortex of monkeys (Tömböl, 1974), the above assumptions would suggest that chandelier cells form about 1.5% of the neuronal population of the monkey visual cortex. This is probably an overestimate.

(Emson & Hunt, 1981; Fuxe, Hökfelt, Said, & Mutt, 1977). These cells can be stained immunohistochemically. About 1350 VIP cells are estimated to occur under 1 mm^2 of surface in the rat visual cortex (this figure is uncorrected for shrinkage) (Morrison, Magistretti, Benoit, & Bloom, 1984), thereby making up about 1% of the total neurons. Whether bipolar VIP cells also have other nonpeptide transmitters is not yet known. The type of synapse made by each type of bipolar cell is not clear. Their relationship, if any, to the "columns" reported in the cortex is also unclear. Nevertheless, their narrow shape is intriguing.

The Behavior of Single Neurons in the Cerebral Cortex

The behavior of single neurons in the cerebral cortex will depend to some extent on the exact nature of the neuron concerned. Here we will treat only the general problem.

It is widely believed that the dendrites of most neurons behave in a largely passive manner. Their cable constants can be estimated from passive membrane transients that can be recorded intracellularly. In the cortex the total soma-dendritic input resistance ranges from 6.7 to 78 megohms with a mean of 24 megohms, while the membrane time constant tends to be relatively constant at around 8.2 milliseconds (Connors, Gutnick, & Prince, 1982). If the specific membrane capacitance is assumed to be about 1 $\mu F/cm^2$ (K. S. Cole, 1968), then the mean specific membrane resistance is relatively high at about 8,200 ohm·cm^2. This implies a relatively large length constant for most neurons in the cerebral cortex.

However, there are disturbing reports that the dendrites of some neurons (for example, the Purkinje cells of the cerebellum and the pyramidal cells of the hippocampus) may have spike generating patches (Llinás & Nicholson, 1969; Spencer & Kandel, 1961). Unfortunately, the experimental techniques to look for this are not easy and the interpretation of the results is not straightforward.

Clearly, it would make a tremendous difference if the dendrites of a neuron were not purely passive. To give just one example, if its dendrites are passive it might be argued that the apical dendrites of a pyramidal cell may mainly serve to "modulate" the inputs of the basal dendrites, because the apical dendritic shaft would attenuate any change of potential produced by synapses in the apical tuft so that, by itself, it might not be able to fire the cell. If the apical dendrite had spike generating capabilities, this argument would not be valid. It is clearly an urgent matter to decide just how potential changes are propagated in dendrites of different types of cells.

Another important theoretical parameter is the "weight" of individual synapses; that is, the size (and time course) of the potential change that synapse produces in the postsynaptic cell. Since in many theories these weights not only determine the behavior of the neuron but are thought to be important for memory, they are obviously of considerable significance. Such weights could be influenced by many factors, both presynaptic and postsynaptic. Moreover, the weights could be subject to transient changes due to many different biochemical processes. Parenthetically we may point out that there is a problem concerning long-term memory and the synaptic weights that may be associated with it. How do synapses manage to remember anything over a period of years in the face of relentless molecular turnover? This makes one wonder whether some single structural feature, either at the molecular level or at a higher level, might embody the "long-term weight." A naive guess might be that it is simply the *area* of the synapse, since in the neocortex this varies, from synapse to synapse, by a factor of 10 (Peters & Kaiserman-Abramof, 1969).

Type II synapses tend to be nearer the axon hillock than do Type I synapses. Thus, it can be argued that such inhibitory synapses on a given dendrite can exercise a partial veto on more distal excitation and thus the exact arrangement of the various synapses could be significant. On this view, a single neuron, rather than being a single integrating device for distributed excitation and inhibition, may be a more complex processing unit, with each dendritic branch acting, in a loose sense, as an integrating unit for its own inputs (Koch, Poggio, & Torre, 1982). It remains to be seen whether this new concept is really valid. Whether it is or not, it should be noted that it is a lot easier to implement addition, subtraction, and division than it is to implement multiplication in a single neuron. In logical terms, AND-NOT seems easier to implement than AND. However, because of the uncertainties in our knowledge, such generalizations are precarious.

In any single neuron a synapse nearer the cell body is likely to have a larger effect than one near the ends of the dendrites, but even this rather obvious deduction has been questioned for neurons whose dendrites generate action potentials. It is obvious that a lot needs to be learned about dendritic behavior before theorists have a solid body of facts to build on.

The Behavior of Groups of Neurons in the Cerebral Cortex

If little is known for certain about single neurons, even less is known about neuronal groups and their behavior in the cerebral cortex.

As we have already seen, in the cortex there are extensive axon collaterals. Many of these only extend laterally for relatively small distances—less than mm—but a significant fraction spread for several millimeters (Gilbert & Wiesel, 1983; Rockland & J. S. Lund, 1983). However, they never seem to spread as much as a centimeter, though this is not true for the diffuse inputs such as those from the brain stem. Thus, in mathematical terms, for one cortical area the connections seem to be "near-diagonal," assuming that we have a two-dimensional arrangement of cells and a four-dimensional connection matrix. Whether excitatory axons spread more or less than inhibitory axons is not clear. The data on this point are confusing.

A favorite theoretical model is one in which all cells of one type connect directly to all other cells of the same type. It seems very rare for a cell to connect to itself (the diagonal term), although occasional axon collaterals that terminate upon their parent cell have been described (Van der Loos & Glaser, 1972). A more disturbing criticism is that, among the sea of axon collaterals, we really have no evidence that cells of one type connect to other cells of the same type. A better guess would be that the collaterals usually contact cells of *other* types in that area of cortex, often to cells in a different layer. Our ignorance springs partly from the fact that we lack convenient experimental methods for studying which cell types actually connect to which.

Are there any general rules about the connection in the neocortex? An outline model, based on recent studies of the visual cortex (Gilbert, 1983; Wiesel & Gilbert, 1983), might be that the main extrinsic input to a cortical area (from the thalamus) is to the middle layer. From there the excitation spreads largely to the upper layer, and from there to the deep layer—first to layer V and then to layer VI. The flow of information for other inputs is less clear. Inputs from other cortical areas tend to end in the upper layer and presumably spread from there to the deep layer. This description is certainly grossly oversimplified in almost every respect but it may turn out to have an element of truth in it.

Rates of Firing

The average discharge rate of neocortical neurons is relatively slow—perhaps only 50 to 100 spikes per second, though the rate for a brief time can rise as high as several hundred spikes per second (see e.g., Lynch, Mountcastle, Talbot, & Yin, 1977). This means that the time between spikes is often 10 to 20 milliseconds. Bearing in mind how much "computation" the brain can do in 100 milliseconds, this

seems to suggest that the *average* rate of firing of an individual neuron can only be transmitted rather approximately and that computations involving many reiterations are unlikely, at least for the initial processing of an input. The precise nature of this limitation deserves further study.

Most, but not all, cortical neurons have very low resting discharges: only a few spikes a second (see, e.g., Abeles, 1982). This is likely to cause problems when one needs to signal both positive and negative values of a function. One obvious way is for one set of neurons to signal the positive values and another to signal the negative ones. In the visual system this appears to be initiated by the retinal ganglion cells (which have separate ON-center and OFF-center receptive fields, Kuffler, 1953) which signal information from the retina to the lateral geniculate body (Cajal, 1892) and from there are relayed to the primary visual cortex and elsewhere.

A large "soma-dendritic" spike generally follows an action potential which is initiated at the axon hillock. This large spike is believed to take place at the cell body and possibly invade the proximal parts of the dendrites (Brock, Coombs, & Eccles, 1952a). This might conceivably wipe clean the slate each time a neuron fires, as well as tell each synapse that the cell had fired. There is evidence suggesting that such a mechanism might exist in spinal motoneurons. Virtually complete destruction of pre-existent postsynaptic potentials have been seen following soma-dendritic spikes in spinal motoneurons (Brock, Coombs, & Eccles, 1952b). Further data on this important point would be welcome.

Feature Detection

In certain parts of the neocortex, especially those near the sensory inputs, it has been shown that a particular neuron will respond best to a certain set of "features" in the input. Thus a so-called "simple cell" in the first visual area responds best to a line or edge of a particular orientation in a particular place in the visual field (Hubel & Wiesel, 1968). It may prefer a certain small range of spatial frequencies (Schiller, Finlay, & Volman, 1976). It may respond better to movement in one direction rather than that in the opposite direction (Hubel & Wiesel, 1968). Some cells are sensitive to the horizontal "disparity" between the input to the two eyes (Hubel & Wiesel, 1970a). Other cells may respond better to some wavelengths of light than to others (Michael, 1978) and so on.

The description and classification of features are matters of the first importance, but even for the early parts of the visual system our knowledge is still very incomplete. These details are of the greatest interest to theorists actively studying a particular system. Here we can only mention them in passing.

As one proceeds further from the sensory input, the mapping of the periphery becomes more diffuse. At the same time most of the neurons tend to respond to one or another feature of the stimulus—movement, color, etc.—while still others respond to more elaborate features in the input. For example, Zeki has shown that in the visual area V4 of the macaque, a cell's response to the wavelength of the light depends somewhat on the wavelength coming from fairly distant regions in the visual field (Zeki, 1978). In the first visual area, on the other hand, a neuron's response to wavelength is much more local (Michael, 1978). This makes sense since although the axon collaterals spread a similar cortical distance in both areas, V4 is smaller than V1, so that in the former the collaterals can reach more distant parts of the visual field, especially as the "mapping" in V4 is more diffuse than it is in V1 (Van Essen & Zeki, 1978).

As one proceeds further into the system it becomes increasingly difficult to discover exactly what "feature" a cell likes best. A few neurons in the cortex, lining the superior temporal sulcus and in the inferior temporal cortex of monkeys, appear to respond to pictures of faces (Bruce, Desimone, & Gross, 1981; Desimone, Albright, Gross, & Bruce, 1980; Perrett, Rolls, & Caan, 1982). If the eyes are blocked out in the pictures, the neurons fire less. If the features of the face are jumbled, the neurons do not respond at all. The exact location and the orientation of the pictures do not appear to be critical to the response of these neurons. It is claimed that other neurons in these complex cortical areas respond specifically to hands (Desimone et al., 1980), but for most of the neurons in these cortical areas, the best type of stimulus has eluded discovery. This is a case where a good theory might give useful pointers to an experimentalist. Since a neuron's output is simply the spikes it sends down its axon, the same pattern of spikes can conceivably be produced by a whole variety of different but related inputs. Thus, in a very real sense the firing of a single neuron conveys somewhat ambiguous information. It is widely believed that it is the pattern of a set of neurons which is best thought of as conveying information. It thus becomes important to know whether the input from a *single* cortical neuron can, by itself, fire a particular cell in its projection field or whether several neurons are required to do this. There is little evidence in the neocortex that a single neuron can, by itself, fire a cell, but exactly how many are needed can only be guessed at. Even where we know that the axon of one neuron contacts the

dendrites of another neuron, we usually do not know *how many* distinct synapses it makes on the second cell. Methods that could answer this question would be very valuable.

Conceptually the problem is not straightforward, since the background "noise" (the occasional random firing of many neurons) is likely to bring a neuron closer to threshold and thus reduce the number of more active neurons required to fire the cell. Spontaneous, transient fluctuations of the membrane potential have been seen in all intracellularly examined cortical neurons (Connors, Gutnick, & Prince, 1982).

There seems to be rather little *correlated* firing of neurons in the neocortex. (For details, see the monograph by Abeles, 1982.) That is, neighboring neurons seldom fire at precisely the same time. When neurons do fire with some degree of synchrony, as they appear to do to produce alpha-waves in the EEG, they do so mainly when the mind appears to be idle. This apparent lack of synchronized firing suggests that the brain is not organized, as a modern digital computer is, on a repeating time cycle. However, the thalamus does appear to impose a degree of rhythm on its output, so it is possible that more subtle time effects exist and have been missed.

An important unresolved question concerning feature detection is whether it is inborn or learned. This is a difficult topic and we can only touch on it. At the moment it seems likely that "feature detectors" are to a considerable extent inborn but can be tuned up by experience, especially during certain "critical periods" in development. The neocortex does not appear in its details to be a general-purpose computing machine. Each area (with its connections) seems designed to carry out a specific function, even though this can be modified somewhat by experience. Natural selection has had every opportunity to develop the necessary neuronal machinery to cope with issues which, because of the nature of the external world, have changed little in evolution. Theorists almost always assume that they are cleverer than natural selection. This is usually a mistake.

In spite of the above remarks, it is reasonable to assume that the neocortex has evolved because it is good at a particular sort of computation and that, with appropriate local variations, it may be broadly the same in all parts. We may also expect that these basic processes will be complicated by somewhat elaborate neural gadgetry designed to make for better performance. What these basic computational processes are remains to be discovered.

Other sections of this book concern models constructed out of "units." These units often have properties similar in some respects to neurons in that they have multiple inputs, some sort of summation rule, a threshold rule, and a single output which is usually distributed to several other units. However, their inventors are always careful to

point out that they are not intended to represent real neurons. Indeed, at this stage in the game, it would be foolish to attempt to do this. Nor are most of the units used satisfactory idealizations of real neurons. If the properties of real neurons present useful gadgets to neural modelers, they should not be mixed together in combinations that never occur together in the brain.

Another explanation offered by modelers in defense of their units is that a single unit really stands for a *group* of neurons. This might be acceptable to neuroscientists if it were carefully stated how this group might be built out of more or less real neurons, but this is seldom if ever done. Consequently, it is difficult to know whether a given "unit" is plausible or not.

Another approach to the difficulty is to claim that "units are placeholders for informational states" (J. L. McClelland, personal communication) and that the relationship between the neurons and such states may be complex. This may indeed be plausible, but from the neuroscientist's point of view it makes it almost impossible to test the models unless the relationship is spelled out in detail.

Another difficulty is that neural modelers seldom state exactly what their models are supposed to demonstrate. This difficult question is addressed more fully in the last chapter of this book.

Meanwhile we list here briefly some of the devices loved by theorists which, if interpreted literally, are not justified by the available experimental evidence:

- Neurons that excite some cells and inhibit others.
- Neurons that merely "change sign." For example, a neuron that accepts excitation from one neuron only and whose onput produces inhibition on one neuron only.
- Neurons that connect to all other cells of the same type.
- Neurons with distinctive synapses that do elaborate computations. Apart from spines which sometimes have both a Type I and a Type II synapse, such complications are rare or absent in the neocortex, though they do occur in the thalamus. However, separate dendrites may perform such a role in the cortex.
- A neuron that, by itself, can fire another cell. This does occur in the cerebellum (the climbing fiber on a Purkinje cell) (Eccles, Llinás, & Sakaski, 1966). It is not certain that it does not occur in the neocortex but the available evidence suggests that it is not common. However, chandelier cells and basket cells may, by themselves, be able to veto the firing of another cell.

The following are features found in the neocortex but often not used in theoretical models:

- Veto cells, which appear to veto many other cells and which probably need the summated activity of several distinct inputs to fire them.

- The various diffuse inputs, from the brain stem and elsewhere, which may be important, not only for the general level of arousal of the cortex (as in sleep) but also for potentiating the synaptic modification involved in laying down a memory.

ACKNOWLEDGMENTS

We thank our colleagues at the Salk Institute for many helpful discussions. We are especially grateful to W. Maxwell Cowan and Jay McClelland for their detailed comments on the manuscript. We also thank Betty Lang for her careful typing of the various drafts.

This work was supported by the J. W. Kieckhefer Foundation, the Samuel Roberts Noble Foundation, the System Development Foundation, and NIH Fellowship NS-07061.