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.

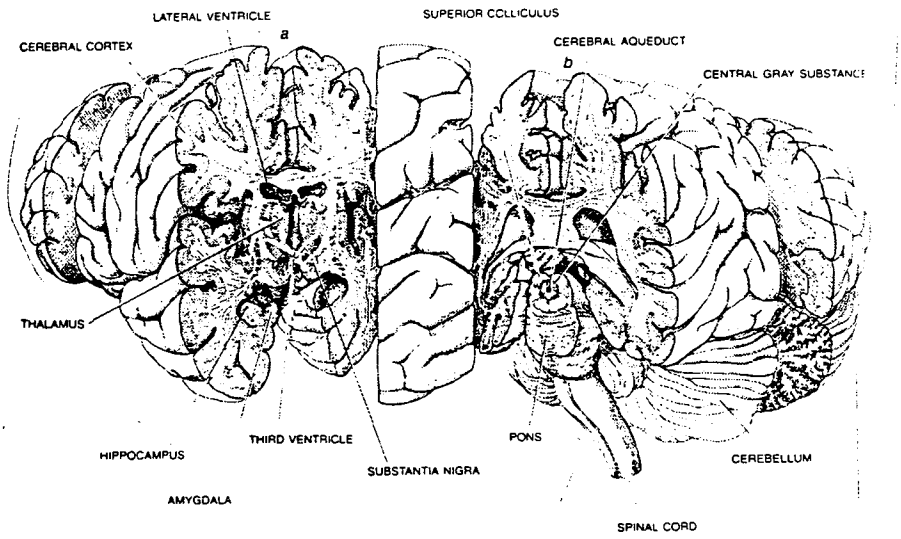


FIGURE 1. The human brain. The cerebral cortex is depicted transparently in this drawing so that some of the internal brain structures are visible. (From "The Organization of the Brain" by W. J. H. Nauta and M. Feirtag, 1979, *Scientific American*, *241*, p. 102. Copyright 1979 by W. H. Freeman & Co.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

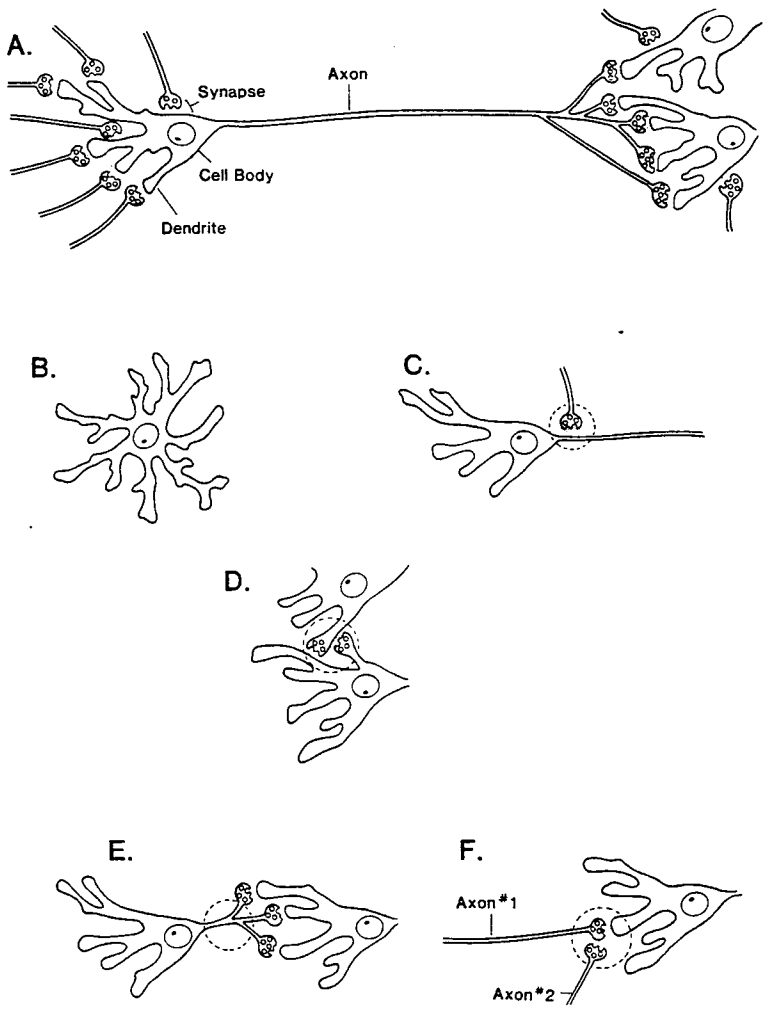


FIGURE 2. Highly schematized diagrams of the "classical" neuronal profile (*A*) and some of its variants (*B-F*). *A*: The "classical" neuron receives synapses on its dendrites and generates action potentials which travel down the axon. The axon subsequently branches and forms synapses on the dendrites 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.

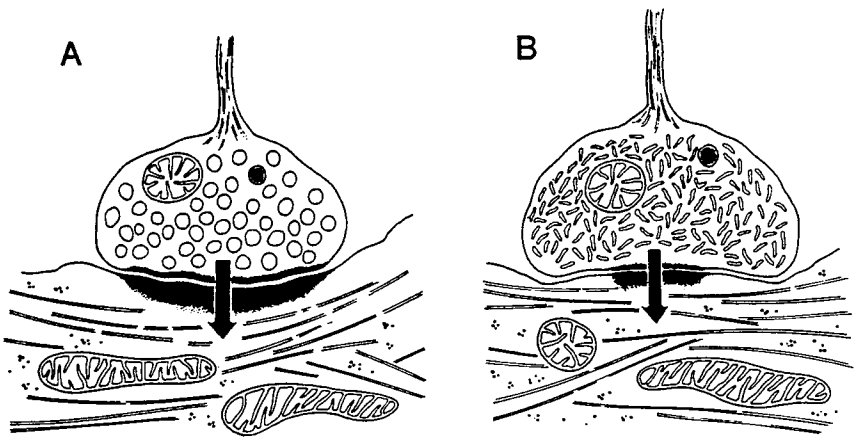


FIGURE 3. Idealized diagrams of a Type I (*A*) and a Type II (*B*) synapse. See text for explanation.

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.

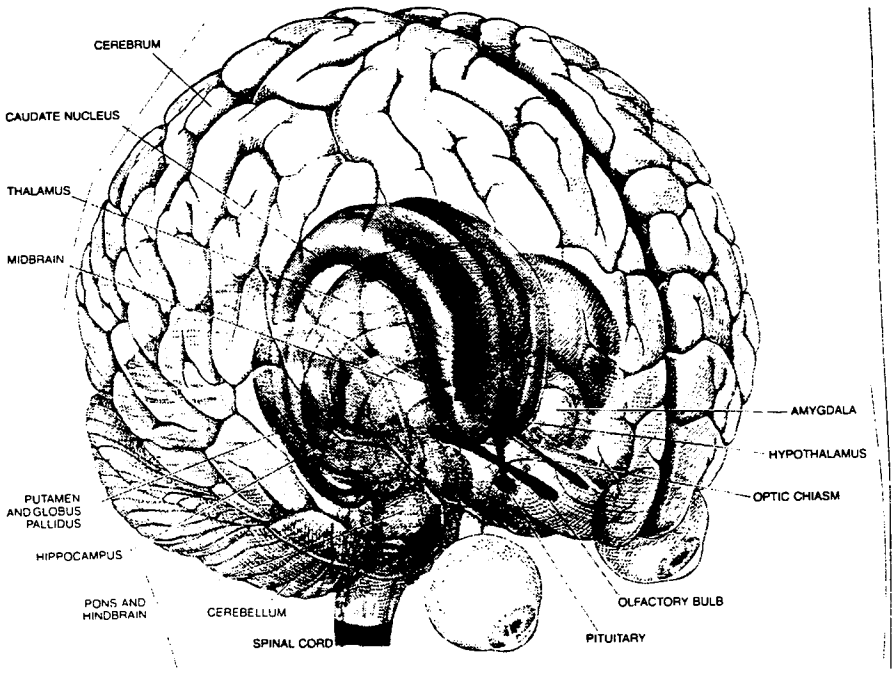


FIGURE 4. A human brain is sliced and opened like a book to demonstrate the continuity of each cortical sheet and its relation to some internal structures. (From "The Organization of the Brain" by W. J. H. Nauta and M. Feirtag, 1979, *Scientific American*, *241*, p. 92. Copyright 1979 by W. H. Freeman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

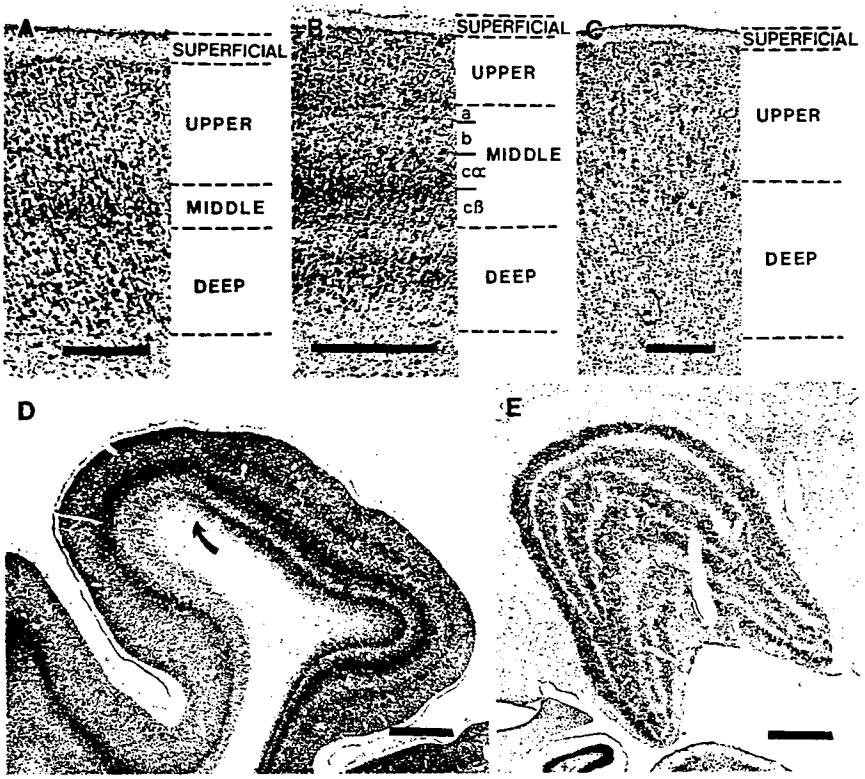


FIGURE 5. *Top*: Some examples of variations in cortical stratification patterns. The stain used is selective for cell bodies. The surface of the brain is at the top in each of these photomicrographs. *A*: Parietal cortex. In most cortical areas, the four main layers are easily recognized. *B*: Striate 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.

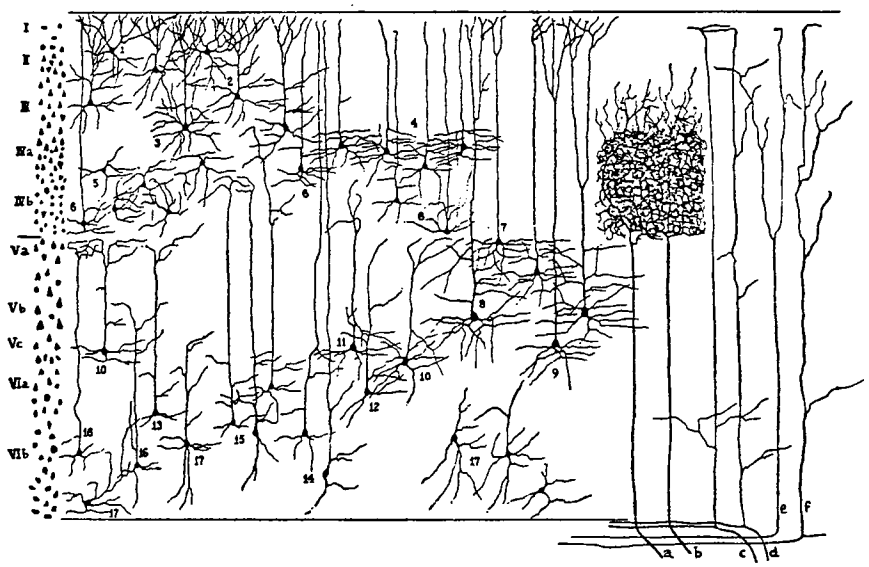


FIGURE 6. The pronounced vertical orientation of many of the dendritic and axonal processes in the neocortex is evident in this diagram of the parietal cortex of an adult mouse. At the left is a diagrammatic representation of all neuronal cell bodies within one very thin section; at the center are the 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.

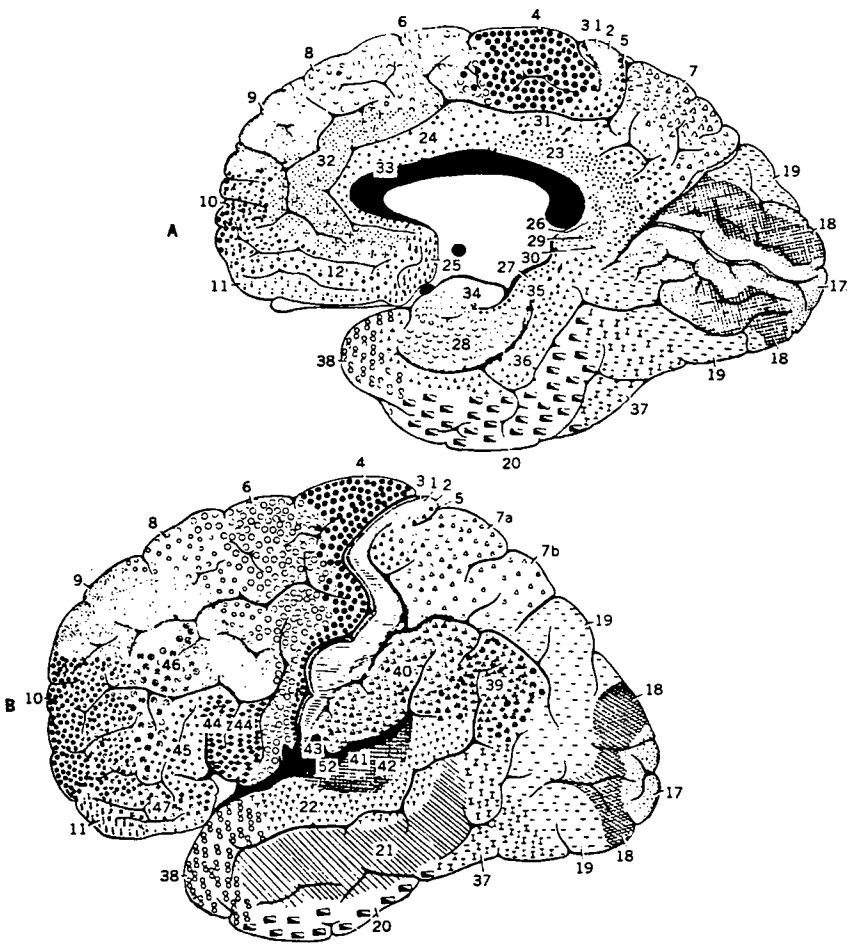


FIGURE 7. Brodmann's areas of the human cerebral cortex. Each of his areas are numbered and indicated by different symbols. *A*: Medial surface of the cerebral cortex (the black areas are occupied by fiber bundles crossing the midline to connect the two hemispheres). *B*: Lateral surface of the cerebral 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

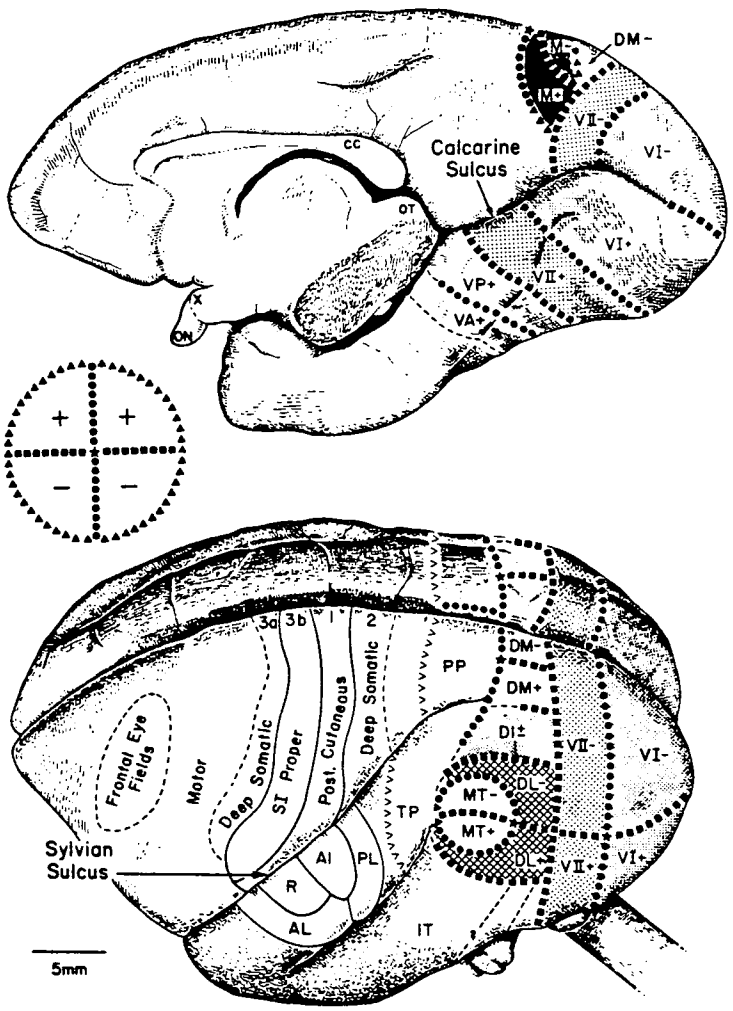


FIGURE 8. Visual processing areas of the owl monkey cerebral cortex. Each of the different types of shading represent different areas. +'s indicate the regions representing the dorsal half of the visual field. −'s indicate the ventral half of the visual field. DI, dorsointermediate visual area; DL, dorsolateral 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

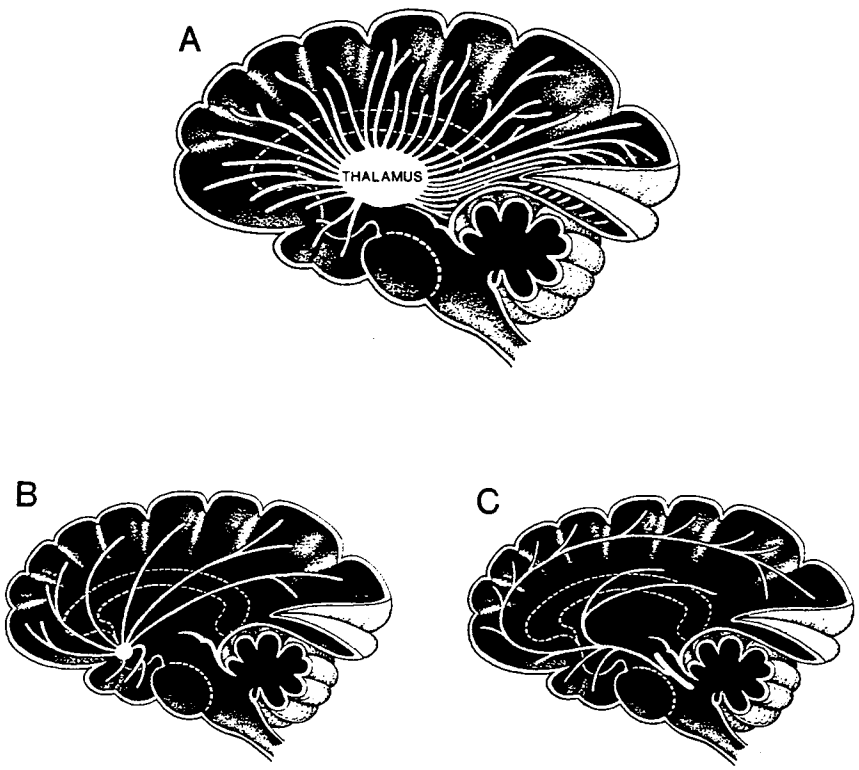


FIGURE 9. Some of the inputs to the neocortex. *A*: Most of the information entering the neocortex gets there through the thalamus. *B*: A diffuse cholinergic input arises in the basal forebrain. *C*: Diffuse noradrenergic and serotonergic inputs arise in the brain stem.

[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).