


2. IMPULSES, SYNAPSES, AND CIRCUITS

A large part of neuroscience concerns the nuts and bolts of the subject: how single cells work and how information is conveyed from cell to cell across synapses. It should be obvious that without such knowledge we are in the position of someone who wants to understand the workings of a radio or TV but does not know anything about resistors, condensers, or transistors. In the last few decades, thanks to the ingenuity of several neurophysiologists, of whom the best known are Andrew Huxley, Alan Hodgkin, Bernard Katz, John Eccles, and Stephen Kuffler, the physicochemical mechanisms of nerve and synaptic transmission have become well understood. It should be equally obvious, however, that this kind of knowledge by itself cannot lead to an understanding of the brain, just as knowledge about resistors, condensers, and transistors alone will not make us understand a radio or TV, or knowledge of the chemistry of ink equip us to understand a Shakespeare play. In this chapter I begin by summing up part of what we know about nerve conduction and synaptic transmission. To grasp the subject adequately, it is a great help to know some physical chemistry and electricity, but I think that anyone can get a reasonable feel for the subject without that. And in any case you only need a very rudimentary understanding of these topics to follow the subsequent chapters. The job of a nerve cell is to take in information from the cells that feed into it, to sum up, or *integrate*, that information, and to deliver the integrated information to other cells. The information is usually conveyed in the form of brief events called *nerve impulses*. In a given cell, one impulse is the same as any other; they are stereotyped events. At any moment a cell's rate of firing impulses is determined by information it has just received from the cells feeding into it, and its firing rate conveys information to the cells that it in turn feeds into. Impulse rates vary from one every few seconds or even slower to about 1000 per second at the extreme upper limit.




This cross-sectional microscopic drawing of the nerve cells in the retina was made by Santiago Ramon y Cajal, the greatest neuroanatomist of all time. From the top, where the slender rods and fatter cones are shown, to the bottom, where optic nerve fibers lead off to the right, the retina measures one-quarter millimeter.

THE MEMBRANE POTENTIAL

What happens when information is transferred from one cell to another at the synapse? In the first cell, an electrical signal, or *impulse*, is initiated on the part of an axon closest to the cell body. The impulse travels down the axon to its terminals. At each terminal, as a result of the impulse, a chemical is released into the narrow, fluid-filled gap between one cell and the next—the *synaptic cleft*—and diffuses across this 0.02-micrometer gap to the second, postsynaptic, cell. There it affects the membrane of the second cell in such a way as to make the second cell either more or less likely to fire impulses. That is quite a mouthful, but let's go back and examine the process in detail. The nerve cell is bathed in and contains salt water. The salt consists not only of sodium chloride, but also of potassium chloride, calcium chloride, and a few less common salts. Because most of the salt molecules are ionized, the fluids both inside and outside the cell will contain chloride, potassium, sodium, and calcium ions (Cl^- , K^+ , Na^+ and Ca^{2+}). In the *resting*

state, the inside and outside of the cell differ in electrical potential by approximately one-tenth of a volt, positive outside. The precise value is more like 0.07 volts, or 70 millivolts. The signals that the nerve conveys consist of transient changes in this resting potential, which travel along the fiber from the cell body to the axon endings. I will begin by describing how the charge across the cell membrane arises. The nerve-cell membrane, which covers the entire neuron, is a structure of extraordinary complexity. It is not continuous, like a rubber balloon or hose, but contains millions of passages through which substances can pass from one side to the other. Some are pores, of various sizes and shapes. These are now known to be proteins in the form of tubes that span the fatty substance of the membrane from one side to the other. Some are more than just pores; they are little machine-like proteins called *pumps*, which can seize ions of one kind and bodily eject them from the cell, while bringing others in from the outside. This pumping requires energy, which the cell ultimately gets by metabolizing glucose and oxygen. Other pores, called *channels*, are valves that can open and close. What influences a given pore to open or close depends on what kind of pore it is. Some are affected by the charge across the membrane; others open or close in response to chemicals floating around in the fluid inside or outside the cell. The charge across the membrane at any instant is determined by the concentrations of the ions inside and out and by whether the various pores are open or closed. (I have already said that pores are affected by the charge, and now I am saying that the charge is determined by the pores. Let's just say for now that the two things can be interdependent. I will explain more soon.) Given the existence of several kinds of pores and several kinds of ions, you can see that the system is complicated. To unravel it, as Hodgkin and Huxley did in 1952, was an immense accomplishment.




A synapse appears as the thin, dark area near the bottom center in this electron microscope picture of a section through cerebellar cortex of a rat. To the left of the synapse, an axon cut in cross section is filled with tiny round synaptic vesicles, in which neurotransmitter is stored. To the right a dendritic process (called a spine) can be seen coming off of a large dendritic branch, which runs horizontally across the picture near the top. (The two sausage-like dark structures in this dendrite are mitochondria.) The two membrane surfaces, of the axon and dendrite, come together at the synapse, where they are thicker and darker. A 20-nanometer cleft separates them.

First, how does the charge get there? Suppose you start with no charge across the

membrane and with the concentrations of all ions equal inside and outside. Now you turn on a pump that ejects one kind of ion, say sodium, and for each ion ejected brings in another kind, say potassium. The pump will not in itself produce any charge across the membrane, because just as many positively charged ions are pumped in as are pumped out (sodium and potassium ions both having one positive charge). But now imagine that for some reason a large number of pores of one type, say the potassium pores, are opened. Potassium ions will start to flow, and the rate of flow through any given open pore will depend on the potassium concentrations: the more ions there are near a pore opening, the more will leak across, and because more potassium ions are inside than outside, more will flow out than in. With more charge leaving than entering, the outside will quickly become positive with respect to the inside. This accumulation of charge across the membrane soon tends to discourage further potassium ions from leaving the cell, because like charges repel one another. Very quickly—before enough K^+ ions cross to produce a measurable change in the potassium concentration—the positive-outside charge builds up to the point at which it just balances the tendency of K^+ ions to leave. (There are more potassium ions just inside the pore opening, but they are repelled by the charge.) From then on, no net charge transfer occurs, and we say the system is in equilibrium. *In short, the opening of potassium pores results in a charge across the membrane, positive outside.* Suppose, instead, we had opened the sodium pores. By repeating the argument, substituting "inside" for "outside", you can easily see that the result would be just the reverse, a negative charge outside. If we had opened both types of pores at the same time, the result would be a compromise. To calculate what the membrane potential is, we have to know the relative concentrations of the two ions and the ratios of open to closed pores for each ion—and then do some algebra.

THE IMPULSE


When the nerve is at rest, most but not all potassium channels are open, and most sodium channels are closed; the charge is consequently positive outside. During an impulse, a large number of sodium pores in a short length of the nerve fiber suddenly open, so that briefly the sodium ions dominate and that part of the nerve suddenly becomes negative outside, relative to inside. The sodium pores then reclose, and meanwhile even more potassium pores have opened than are open in the resting state. Both events—the sodium pores reclosing and additional potassium pores opening—lead to the rapid restoration of the positive-outside resting state. The whole sequence lasts about one-thousandth of a second. All this depends on the circumstances that influence pores to open and close. For both Na^+ and K^+ channels, the pores are sensitive to the charge across the membrane.



Top: A segment of nerve axon at rest. The sodium pump has expelled most sodium ions and brought in potassium ions. Sodium channels are mainly closed. Because many potassium channels are open, enough potassium ions have left relative to those entering to charge the membrane to 70 millivolts positive outside. *Bottom:* A nerve impulse is traveling from left to right. At the extreme right the axon is still in the resting state. In the middle section the impulse is in full swing: sodium channels are open, sodium ions are pouring in (though not in nearly large enough amounts to produce any measurable changes in concentration in the course of one impulse); the membrane is now 40 millivolts, negative outside. At the extreme left the membrane is recovering. The resting potential is restored, because more potassium channels have opened (and then closed) and because sodium channels have automatically closed. Because sodium channels cannot immediately reopen, a second impulse cannot occur for about a millisecond. This explains why the impulse when under way cannot travel back toward the cell body.

Making the membrane less positive outside—*depolarizing* it from its resting state—results in the opening of the pores. The effects are not identical for the two kinds of pores: the sodium pores, once opened, close of their own accord, even though the depolarization is maintained, and are then incapable of reopening for a few thousandths of a second; the potassium pores stay open as long as the depolarization is kept up. For a given depolarization, the number of sodium ions entering is at first greater than the number of potassium ions leaving, and the membrane swings negative outside with respect to inside; later, potassium dominates and the resting potential is restored. In this sequence of events constituting an impulse, in which pores open, ions cross, and the membrane potential changes and changes back, the number of ions that actually cross the membrane—sodium entering and potassium leaving—is miniscule, not nearly enough to produce a measurable change in the concentrations of ions inside or outside the cell. In several minutes a nerve might fire a thousand times, however, and that might be enough to change the concentrations, were it not that the pump is meanwhile continually ejecting sodium and bringing in potassium so as to keep the concentrations at their proper resting levels. The reason that during an impulse such small charge transfers result in such large potential swings is a simple matter of electricity: the capacitance of the membrane is low, and potential is equal to charge transferred divided by capacitance. A depolarization of the membrane—making it less positive-outside than it is at rest—is what starts up the impulse in the first place. If, for example, we suddenly insert some sodium ions into the

resting fiber, causing a small initial depolarization, a few sodium pores open as a consequence of that depolarization but because many potassium pores are already open, enough potassium like the current in a copper wire. No electricity or ions or anything tangible travels along the nerve, just as nothing travels from handle to point when a pair of scissors closes. (Ions do flow in and out, just as the blades of the scissors move up and down.) It is the event, the intersection of the blades of the scissors or the impulse in the nerve, that travels. Because it takes some time before sodium channels are ready for another opening and closing, the highest rate at which a cell or axon can fire impulses is about 800 per second. Such high rates are unusual, and the rate of firing of a very active nerve fiber is usually more like 100 or 200 impulses per second. One important feature of a nerve impulse is its *all-or-none* quality. If the original depolarization is sufficient—if it exceeds some threshold value, (going from the resting level of 70 millivolts to 40 millivolts, positive outside)—the process becomes regenerative, and reversal occurs all the way to 40 millivolts, negative outside. The magnitude of the reversed potential traveling down the nerve (that is, the impulse) is determined by the nerve itself, not by the intensity of the depolarization that originally sets it going.



The membrane of a glial cell is wrapped around and around an axon, shown in cross section in this enlarged electron microscopic view. The encircling membrane is myelin, which speeds nerve impulses by raising the resistance and lowering the capacitance between inside and outside. The axon contains a few organelles called microtubules.

It is analogous to any explosive event. How fast the bullet travels has nothing to do with how hard you pull the trigger. For many brain functions the speed of the impulse seems to be very important, and the nervous system has evolved a special mechanism for increasing it. Glial cells wrap their plasma membrane around and around the axon like a jelly roll, forming a sheath that greatly increases the effective thickness of the nerve membrane. This added thickness reduces the membrane's capacitance, and hence the amount of charge required to depolarize the nerve. The layered substance, rich in fatty material, is called *myelin*. The sheath is interrupted every few millimeters, at *nodes of Ranvier*, to allow the currents associated with the impulse to enter or leave the axon. The

result is that the nerve impulse in effect jumps from one node to the next rather than traveling continuously along the membrane, which produces a great increase in conduction velocity. The fibers making up most of the large, prominent cables in the brain are myelinated, giving them a glistening white appearance on freshly cut sections. *White matter* in the brain and spinal cord consists of myelinated axons but no nerve cell bodies, dendrites, or synapses. *Grey matter* is made up mainly of cell bodies, dendrites, axon terminals, and synapses, but may contain myelinated axons. The main gaps remaining in our understanding of the impulse, and also the main areas of present-day research on the subject, have to do with the structure and function of the protein channels.

SYNAPTIC TRANSMISSION

How are impulses started up in the first place, and what happens at the far end, when an impulse reaches the end of an axon?

The part of the cell membrane at the terminal of an axon, which forms the first half of the synapse (the presynaptic membrane), is a specialized and remarkable machine. First, it contains special channels that respond to depolarization by opening and letting positively charged *calcium* ions through. Since the concentration of calcium (like that of sodium) is higher outside the cell than inside, opening the gates lets calcium flow in. In some way still not understood, this arrival of calcium inside the cell leads to the expulsion, across the membrane from inside to outside, of packages of special chemicals called neurotransmitters. About twenty transmitter chemicals have been identified, and to judge from the rate of new discoveries the total number may exceed fifty. Transmitter molecules are much smaller than protein molecules but are generally larger than sodium or calcium ions. Acetylcholine and noradrenaline are examples of neurotransmitters. When these molecules are released from the presynaptic terminal they quickly diffuse across the 0.02-micrometer synaptic gap to the postsynaptic membrane. The postsynaptic membrane is likewise specialized: embedded in it are protein pores called *receptors*, which respond to the neurotransmitter by causing channels to open, allowing one or more species of ions to pass through. Just which ions (sodium, potassium, chloride) are allowed to pass determines whether the postsynaptic cell is itself depolarized or is stabilized and prevented from depolarizing. To sum up so far, a nerve impulse arrives at the axon terminal and causes special neurotransmitter molecules to be released. These neurotransmitters act on the postsynaptic membrane either to lower its membrane potential or to keep its membrane potential from being lowered. If the membrane potential is lowered, the frequency of firing increases; we call such a synapse excitatory. If instead the membrane is stabilized at a value above threshold, impulses do not occur or occur less often; in this case, the synapse is termed *inhibitory*. Whether a given synapse is excitatory or inhibitory depends on which neurotransmitter is released and which receptor molecules are present. Acetylcholine, the best-known transmitter, is in some synapses excitatory and in others inhibitory: it excites limb and trunk muscles but inhibits the heart. Noradrenaline is usually excitatory; gamma-amino butyric acid (GABA) is usually inhibitory. As far as we know, a given synapse remains either excitatory or inhibitory for the life of the animal. Any one nerve cell is contacted along its dendrites and cell body by tens, hundreds, or thousands of terminals; at any instant it is thus being told by some synapses to depolarize and by others not to. An impulse coming in over an excitatory

terminal will depolarize the postsynaptic cell; if an impulse comes in simultaneously over an inhibitory terminal, the effects of the two will tend to cancel each other. At any given time the level of the membrane potential is the result of all the excitatory and inhibitory influences added together. A single impulse coming into one axon terminal generally has only a miniscule effect on the next cell, and the effect lasts only a few milliseconds before it dies out. When impulses arrive at a cell from several other nerve cells, the nerve cell sums up, or integrates, their effects. If the membrane potential is sufficiently reduced—if the excitatory events occur in enough terminals and at a high enough rate—the depolarization will be enough to generate impulses, usually in the form of a repetitive train. The site of impulse initiation is usually where the axon leaves the cell body, because this happens to be where a depolarization of a given size is most likely to produce a regenerative impulse, perhaps owing to an especially high concentration of sodium channels in the membrane. The more the membrane is depolarized at this point, the greater the number of impulses initiated every second. Almost all cells in the nervous system receive inputs from more than one other cell. This is called *convergence*. Almost all cells have axons that split many times and supply a large number of other nerve cells—perhaps hundreds or thousands. We call this *divergence*. You can easily see that without convergence and divergence the nervous system would not be worth much: an excitatory synapse that slavishly passed every impulse along to the next cell would serve no function, and an inhibitory synapse that provided the only input to a cell would have nothing to inhibit, unless the postsynaptic cell had some special mechanism to cause it to fire spontaneously. I should make a final comment about the signals that nerve fibers transmit.

Although most axons carry all-or-none impulses, some exceptions exist. If local depolarization of a nerve is subthreshold—that is, if it is insufficient to start up an explosive, all-or-none propagated impulse—it will nevertheless tend to spread along the fiber, declining with time and with distance from the place where it began. (In a propagated nerve impulse, this local spread is what brings the potential in the next, resting section of nerve membrane to the threshold level of depolarization, at which regeneration occurs.) Some axons are so short that no propagated impulse is needed; by passive spread, depolarization at the cell body or dendrites can produce enough depolarization at the synaptic terminals to cause a release of transmitter. In mammals, the cases in which information is known to be transmitted without impulses are few but important. In our retinas, two or three of the five nerve-cell types function without impulses. An important way in which these passively conducted signals differ from impulses—besides their small and progressively diminishing amplitude—is that their size varies depending on the strength of the stimulus. They are therefore often referred to as *graded* signals. The bigger the signal, the more depolarization at the terminals, and the more transmitter released. You will remember that impulses, on the contrary, do not increase in size as the stimulus increases; instead, their repetition rate increases. And the faster an impulse fires, the more transmitter is released at the terminals. So the final result is not very different. It is popular to say that graded potentials represent an example of analog signals, and that impulse conduction, being all or none, is digital. I find this misleading, because the exact position of each impulse in a train is not in most cases of any significance. What matters is the average rate in a given time interval, not the fine details. Both kinds of signals are thus essentially analog.

A TYPICAL NEURAL PATHWAY


Now that we know something about impulses, synapses, excitation, and inhibition, we can begin to ask how nerve cells are assembled into larger structures. We can think of the central nervous system—the brain and spinal cord—as consisting of a box with an input and an output. The input exerts its effects on special nerve cells called *receptors*, cells modified to respond to what we can loosely term "outside information" rather than to synaptic inputs from other nerve cells. This information can take the form of light to our eyes; of mechanical deformation to our skin, eardrums, or semicircular canals; or of chemicals, as in our sense of smell or taste. In all these cases, the effect of the stimulus is to produce in the receptors an electrical signal and consequently a modification in the rate of neurotransmitter release at their axon terminals. (You should not be confused by the double meaning of *receptor*; it initially meant a cell specialized to react to sensory stimuli but was later applied also to protein molecules specialized to react to neurotransmitters.)



This scanning electron microscope picture shows a neuromuscular junction in a frog. The slender nerve fiber curls down over two muscle fibers, with the synapse at the lower left of the picture.

At the other end of the nervous system we have the output: the motor neurons, nerves that are exceptional in that their axons end not on other nerve cells but on muscle cells. All the output of our nervous system takes the form of muscle contractions, with the minor exception of nerves that end on gland cells. This is the way, indeed the only way, we can exert an influence on our environment. Eliminate an animal's muscles and you cut it off completely from the rest of the world; equally, eliminate the input and you cut off all outside influences, again virtually converting the animal into a vegetable. An animal is, by one possible definition, an organism that reacts to outside events and that influences the outside world by its actions. The central nervous system, lying between input cells and output cells, is the machinery that allows us to perceive, react, and remember—and it must be responsible, in the end, for our consciousness, consciences, and souls. One of the

main goals in neurobiology is to learn what takes place along the way—how the information arriving at a certain group of cells is transformed and then sent on, and how the transformations make sense in terms of the successful functioning of the animal.




Many parts of the central nervous system are organized in successive platelike stages. A cell in one stage receives many excitatory and inhibitory inputs from the previous stage and sends outputs to many cells at the next stage. The primary input to the nervous system is from receptors in the eyes, ears, skin, and so on, which translate outside information such as light, heat, or sound into electrical nerve signals. The output is contraction of muscles or secretions from gland cells.

Although the wiring diagrams for the many subdivisions of the central nervous system vary greatly in detail, most tend to be based on the relatively simple general plan schematized in the diagram on this page. The diagram is a caricature, not to be taken literally, and subject to qualifications that I will soon discuss. On the left of the figure I show the receptors, an array of information-transducing nerves each subserving one kind of sensation such as touch, vibration, or light. We can think of these receptors as the first stage in some sensory pathway. Fibers from the receptors make synaptic contacts with a second array of nerve cells, the second stage in our diagram; these in turn make contact with a third stage, and so on. "Stage" is not a technical or widely applied neuroanatomical term, but we will find it useful. Sometimes three or four of these stages are assembled together in a larger unit, which I will call a structure, for want of any better or widely accepted term. These structures are the aggregations of cells, usually plates or globs, that I mentioned in Chapter 1. When a structure is a plate, each of the stages forming it may be a discrete layer of cells in the plate. A good example is the retina, which has three layers of cells and, loosely speaking, three stages. When several stages are grouped to form a larger structure, the nerve fibers entering from the previous structure and those leaving to go to the next are generally grouped together into bundles, called *tracts*. You will notice in the diagram how common divergence and convergence are: how almost as a rule the axon from a cell in a given stage splits on arriving at the next stage and ends on several or many cells, and conversely, a cell at any stage except the first receives synaptic inputs from a few or many cells in the previous stage. We obviously need to amend and qualify this simplified diagram, but at least we have a model to qualify. We must first recognize that at the input end we have not just one but many sensory systems—vision, touch, taste, smell, and hearing—and that each system has its own sets of stages in the

brain. When and where in the brain the various sets of stages are brought together, if indeed they are brought together, is still not clear. In tracing one system such as the visual or auditory from the receptors further into the brain, we may find that it splits into separate subdivisions. In the case of vision, these subsystems might deal separately with eye movements, pupillary constriction, form, movement, depth, or color. Thus the whole system diverges into separate subpathways. Moreover, the subpaths may be many, and may differ widely in their lengths. On a gross scale, some paths have many structures along the way and others few. At a finer level, an axon from one stage may not go to the next stage in the series but instead may skip that stage and even the next; it may go all the way to the motor neuron. (You can think of the skipping of stages in neuroanatomy as analogous to what can happen in genealogy. The present English sovereign is not related to William the Conqueror by a unique number of generations: the number of "greats" modifying the grandfather is indeterminate because of intermarriage between nephews and aunts and even more questionable events.) When the path from input to output is very short, we call it a *reflex*. In the visual system, the constriction of the pupil in response to light is an example of a reflex, in which the number of synapses is probably about six. In the most extreme case, the axon from a receptor ends directly on a motor neuron, so that we have, from input to output, only three cells: receptor, motor neuron, and muscle fiber, and just two synapses, in what we call a *monosynaptic reflex arc*. (Perhaps the person who coined the term did not consider the nerve-muscle junction a real synapse, or could not count to two.) That short path is activated when the doctor taps your knee with a hammer and your knee jumps. John Nicholls used to tell his classes at Harvard Medical School that there are two reasons for testing this reflex: to stall for time, and to see if you have syphilis. At the output end, we find not only various sets of body muscles that we can voluntarily control, in the trunk, limbs, eyes, and tongue, but also sets that subserve the less voluntary or involuntary housekeeping functions, such as making our stomachs churn, our water pass or bowels move, and our sphincters (between these events) hold orifices closed. We also need to qualify our model with respect to direction of information flow. The prevailing direction in our diagram on page 10 is obviously from left to right, from input to output, but in almost every case in which information is transferred from one stage to the next, reciprocal connections feed information back from the second stage to the first. (We can sometimes guess what such feedback might be useful for, but in almost no case do we have incisive understanding.) Finally, even within a given stage we often find a rich network of connections between neighboring cells of the same order. Thus to say that a structure contains a specific number of stages is almost always an oversimplification. When I began working in neurology in the early 1950s, this basic plan of the nervous system was well understood. But in those days no one had any clear idea how to interpret this bucket-brigade-like handing on of information from one stage to the next. Today we know far more about the ways in which the information is transformed in some parts of the brain; in other parts we still know almost nothing. The remaining chapters of this book are devoted to the visual system, the one we understand best today. I will next try to give a preview of a few of the things we know about that system.


We can now adapt our earlier diagram on page 10 to fit the special case of the visual pathway. As shown in the illustration on this page, the receptors and the next two stages are contained in the retina. The receptors are the rods and cones; the optic nerve, carrying the retina's entire output, is a bundle of axons of the third-stage retinal cells, called retinal ganglion cells. Between the receptors and the ganglion cells are intermediate cells, the most important of which are the *bipolar cells*. The optic nerve proceeds to a way station deep in the brain, the lateral geniculate body. After only one set of synapses, the lateral geniculate sends its output to the striate cortex, which contains three or four stages. You can think of each of the columns in the diagram above as a plate of cells in cross section. For example, if we were to stand at the right of the page and look to the left, we would see all the cells in a layer in face-on view. Each of the columns of cells in the figure represents a two-dimensional array of cells, as shown for the rods and cones in the diagram on the next page.



The initial stages of the mammalian visual system have the platelike organization often found in the central nervous system. The first three stages are housed in the retina; the remainder are in the brain: in the lateral geniculate bodies and the stages beyond in the cortex.

To speak, as I do here, of separate stages immediately raises our problem with genealogy. In the retina, as we will see in Chapter 3, the minimum number of stages between receptors and the output is certainly three, but because of two other kinds of cells, some information takes a more diverted course, with four or five stages from input to output. For the sake of convenience, the diagram ignores these detours despite their importance, and makes the wiring look simpler than it really is. When I speak of the retinal ganglion cells as "stage 3 or 4", it's not that I have forgotten how many there are. To appreciate the kind of transfer of information that takes place in a network of this kind, we may begin by considering the behavior of a single retinal ganglion cell. We know from its anatomy that such a cell gets input from many bipolar cells—perhaps 12,100, or 1000—and that each of these cells is in turn fed by a similar number of receptors. As a general rule, all

the cells feeding into a single cell at a given stage, such as the bipolar cells that feed into a single retinal ganglion cell, are grouped closely together. In the case of the retina, the cells connected to any one cell at the next stage occupy an area 1 to 2 millimeters in diameter; they are certainly not peppered all over the retina. Another way of putting this is that none of the connections within the retina are longer than about 1 to 2 millimeters. If we had a detailed description of all the connections in such a structure and knew enough about the cellular physiology—for example, which connections were excitatory and which inhibitory—we should in principle be able to deduce the nature of the operation on the information. In the case of the retina and the cortex, the knowledge available is nowhere near what we require. So far, the most efficient way to tackle the problem has been to record from the cells with microelectrodes and compare their inputs and outputs. In the visual system, this amounts to asking what happens in a cell such as a retinal ganglion cell or a cortical cell when the eye is exposed to a visual image. In attempting to activate a stage-3 (retinal ganglion) cell by light, our first instinct probably would be to illuminate all the rods and cones feeding in, by shining a bright light into the eye. This is certainly what most people would have guessed in the late 1940s, when physiologists were just beginning to be aware of synaptic inhibition, and no one realized that inhibitory synapses are about as plentiful as excitatory ones. Because of inhibition, the outcome of any stimulation depends critically on exactly where the light falls and on which connections are inhibitory and which excitatory. If we want to activate the ganglion cell powerfully, stimulating all the rods and cones that are connected to it is just about the worst thing we can do. The usual consequence of stimulating with a large spot of light or, in the extreme, of bathing the retina with diffuse light, is that the cell's firing is neither speeded up nor slowed down—in short, nothing results: the cell just keeps firing at its own resting rate of about five to ten impulses per second. To increase the firing rate, we have to illuminate some particular subpopulation of the receptors, namely the ones connected to the cell (through bipolar cells) in such a way that their effects are excitatory.

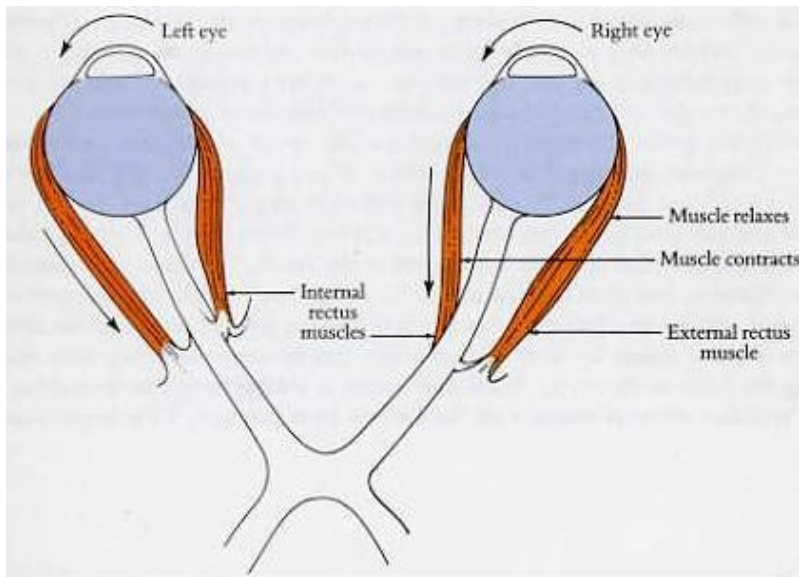


Any one stage in the diagrams on page 10 and on this page consists of a two-dimensional plate of cells. In any one stage the cells may be so densely packed that they come to lie several cells deep; they nevertheless still belong to the same stage.

Illuminating only one such receptor may have hardly any detectable effect, but if we could illuminate all the receptors with excitatory effects, we could reasonably expect their summated influences to activate the cell—and in fact they do. As we will see, for most retinal ganglion cells the best stimulus turns out to be a small spot of light of just the right size, shining in just the right place. Among other things, this tells you how important a role inhibition plays in retinal function.

VOLUNTARY MOVEMENT

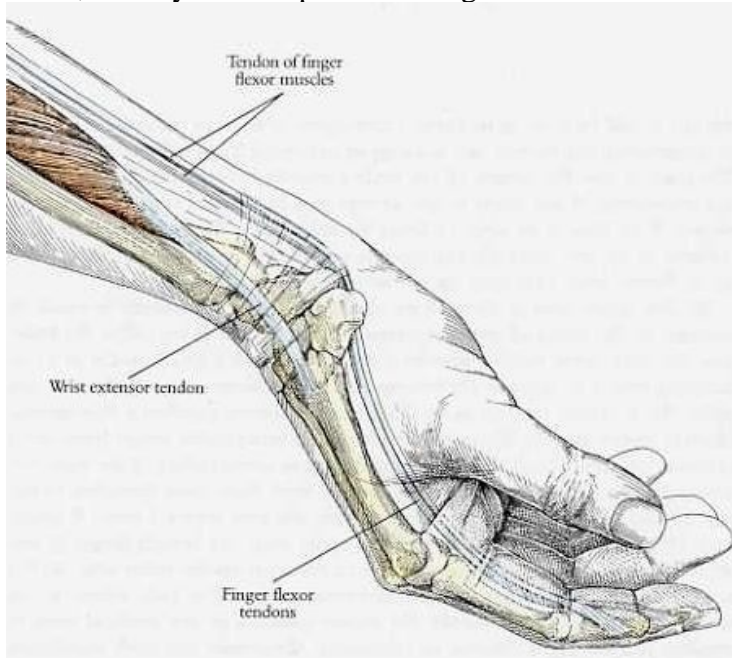
Although this book will concentrate on the initial, sensory stages in the nervous system, I want to mention two examples of movement, just to convey an idea of what the final stages in the diagram on page 10 may be doing. Consider first how our eyes move. Each eye is roughly a sphere, free to move like a ball in a socket. (If the eye did not have to move it might well have evolved as a box, like an old-fashioned box camera.) Each eye has six *extraocular* muscles attached to it and moves because the appropriate ones shorten.



Each eye has its position controlled by six separate muscles, two of which are shown here. These, the external and internal recti, control the horizontal rotation of the eyes, in looking from left to right or from close to far. The other eight muscles, four for each eye, control elevation and depression, and rotation about an axis that in the diagram is vertical, in the plane of the paper.

How these muscles all attach to the eye is not important to us here, but we can easily see from the illustration that for one eye, say the right, to turn inward toward the nose, a person must relax the external rectus and contract the internal rectus muscles. If each muscle did not have some steady pull, or tone, the eye would be loose in its socket; consequently any eye movement is made by contracting one muscle and relaxing its opponent by just the same amount. The same is true for almost all the body's muscle movements. Furthermore, any movement of one eye is almost always part of a bigger complex of movements. If we look at an object a short distance away, the two eyes turn

in; if we look to the left, the right eye turns in and the left eye turns out; if we look up or down, both eyes turn up or down together.



When we flex our fingers by making a fist, the muscles responsible have to pass in front of the wrist and so tend to contract that joint too. The extensors of the wrist have to contract to offset this tendency and keep the wrist stiff.

All this movement is directed by the brain. Each eye muscle is made to contract by the firing of motor neurons in a part of the brain called the *brainstem*. To each of the twelve muscles there corresponds a small cluster of a few hundred motor neurons in the brainstem. These clusters are called *oculomotor nuclei*. Each motor neuron in an oculomotor nucleus supplies a few muscle fibers in an eye muscle. These motor neurons in turn receive inputs from other excitatory fibers. To obtain a movement such as convergence of the eyes, we would like to have these antecedent nerves send their axon branches to the appropriate motor neurons, those supplying the two internal recti. A single such antecedent cell could have its axon split, with one branch going to one oculomotor nucleus and the other to its counterpart on the other side. At the same time we need to have another antecedent nerve cell or cells, whose axons have inhibitory endings, supply the motor neurons to the external recti to produce just the right amount of relaxation. We would like both antecedent sets of cells to fire together, to produce the contraction and relaxation simultaneously, and for that we could have one master cell or group of cells, at still another stage *back* in the nervous system, excite both groups. This is one way in which we can get coordinated movements involving many muscles. Practically every movement we make is the result of many muscles contracting together and many others relaxing. If you make a fist, the muscles in the front of your forearm (on the palm side of the hand) contract, as you can feel if you put your other hand on your forearm. (Most people probably think that the muscles that flex the fingers are in the hand. The hand does contain some muscles, but they happen not to be finger flexors.) As

the diagram on the previous page shows, the forearm muscles that flex the fingers attach to the three bones of each finger by long tendons that can be seen threading their way along the front of the wrist. What may come as a surprise is that in making a fist, you also contract muscles on the back of your forearm. That might seem quite unnecessary until you realize that in making a fist you want to keep your wrist stiff and in midposition: if you flexed only the finger flexor muscles, their tendons, passing in front of the wrist, would flex it too. You have to offset this tendency to unwanted wrist flexion by contracting the muscles that cock back the wrist, and these are in the back of the forearm. The point is that you do it but are unaware of it. Moreover, you don't learn to do it by attending 9 A.M. lectures or paying a coach. A newborn baby will grasp your finger and hold on tight, making a perfect fist, with no coaching or lecturing. We presumably have some executive-type cells in our spinal cords that send excitatory branches both to finger flexors and to wrist extensors and whose function is to subserve fist making. Presumably these cells are wired up completely before birth, as are the cells that allow us to turn our eyes in to look at close objects, without thinking about it or having to learn.