


4. THE PRIMARY VISUAL CORTEX

After Kuffler's first paper on center-surround retinal ganglion cells was published in 1952, the next steps were clear. To account for the properties of the cells, more work was needed at the retinal level. But we also needed to record from the next stages in the visual pathway, to find out how the brain interpreted the information from the eyes. Both projects faced formidable difficulties. In the case of the brain, some years were required to develop the techniques necessary to record from a single cell and observe its activity for many hours. It was even harder to learn how to influence that activity by visual stimulation.




The visual cortex in a monkey, stained by the Golgi method, shows a few pyramidal cells—a tiny fraction of the total number in such a section. The entire height of the photograph represents about 1 millimeter. A tungsten microelectrode, typical of what is used for extra cellular recordings, has been superimposed, to the same scale.

TOPOGRAPHIC REPRESENTATION


Even before further research became possible, we were not completely ignorant about the parts of the brain involved in vision: the geography of the preliminary stages was already well mapped out (see the illustration on the next page). We knew that the optic-nerve fibers make synapses with cells in the lateral geniculate body and that the axons of lateral geniculate cells terminate in the primary visual cortex. It was also clear that these connections, from the eyes to the lateral geniculates and from the geniculates to the cortex, are *topographically organized*. By *topographic representation*, we mean that the mapping of each structure to the next is systematic: as you move along the retina from one point to another, the corresponding points in the lateral geniculate body or cortex trace a continuous path. For example, the optic nerve fibers from a

given small part of the retina all go to a particular small part of the lateral geniculate, and fibers from a given region of the geniculate all go to a particular region of the primary visual cortex. Such an organization is not surprising if we recall the caricature of the nervous system shown in the figure on page 6 in Chapter 1, in which cells are grouped in platelike arrays, with the plates stacked so that a cell at any particular stage gets its input from an aggregate of cells in the immediately preceding stage.



The visual pathway, from eyes to primary visual cortex, of a human brain, as seen from below. Information comes to the two purple-colored halves of the retinas (the right halves, because the brain is seen upside down) from the opposite half of the environment (the left visual field) and ends up in the right (purple) half of the brain. This happens because about half the optic-nerve fibers cross at the chiasm, and the rest stay uncrossed. Hence the rules: each hemisphere gets input from both eyes; a given hemisphere gets information from the opposite half of the visual world.

In the retina, the successive stages are in apposition, like playing cards stacked one on top of the other, so that the fibers can take a very direct route from one stage to the next. In the lateral geniculate body, the cells are obviously separated from the retina, just as, equally obviously, the cortex is in a different place from the geniculate. The style of connectivity nevertheless remains the same, with one region projecting to the next as though the successive plates were still superimposed. The optic-nerve fibers simply gather into a bundle as they leave the eye, and when they reach the geniculate, they fan out and end in a topographically orderly way. (Oddly, between the retina and geniculate, in the optic nerve, they become almost completely scrambled, but they sort out again as they reach the geniculate.) Fibers leaving the geniculate similarly fan out into a broad band that extends back through the interior of the brain and ends in an equally orderly way in the primary visual cortex. After several synapses, when fibers leave the primary visual cortex and project to several other cortical regions, the topographic order is again preserved. Because convergence occurs at every stage, receptive fields tend to become larger: the farther along the path we go, the more fuzzy this representation-by-mapping of the outside world becomes. An important, long-recognized piece of evidence that the pathway is topographically organized comes from clinical observation. If you damage a certain part of your primary visual cortex, you develop a local blindness, as though you had destroyed the corresponding part of your retina. The visual world is thus systematically mapped onto the geniculate and cortex. What was not at all clear in the 1950s was what the mapping might mean. In those days it was not obvious that the brain operates on the information it receives, transforming it in such a way as to make it more useful. People had the feeling that the visual scene had made it to the brain; now the problem for the brain was to make sense of it—or perhaps it was not the brain's problem, but the mind's. The message of the next chapters will be that a structure such as the primary visual cortex does exert profound transformations on the information it receives. We still know very little about what goes on beyond this stage, and in that sense you might argue that we are not much better off. But knowing that one part of the cortex works in a rational, easily understood way gives grounds for optimism that other areas will too. Some day we may not need the word *mind* at all.



A microscopic cross-sectional view of the optic nerve where it leaves the eye, interrupting the retinal layers shown at the left and right. The full width of the picture is about 2 millimeters. The clear area at the top is the inside of the eye. The retinal layers, from the top down, are optic-nerve fibers (clear), the three stained layers of cells, and the black layer of melanin pigment.


RESPONSES OF LATERAL GENICULATE CELLS

The fibers coming to the brain from each eye pass uninterrupted through the *optic chiasm* (from chi, X, the Greek letter whose shape is a cross). There, about half the fibers cross to the side of the brain opposite the eye of origin, and half stay on the same side. From the chiasm the fibers continue to several destinations in the brain. Some go to structures that have to do with such specific functions as eye movements and the pupillary light reflex, but most terminate in the two lateral geniculate bodies. Compared with the cerebral cortex or with many other parts of the brain, the lateral geniculates are simple structures: all or almost all of the roughly one and one half million cells in each geniculate nucleus receive input directly from optic-nerve fibers, and most (not all) of the cells send axons on to the cerebral cortex. In this sense, the lateral geniculate bodies contain only one synaptic stage, but it would be a mistake to think of them as mere relay stations. They receive fibers not only from the optic nerves but also back from the cerebral cortex, to which they project, and from the brainstem reticular formation, which plays some role in attention or arousal. Some geniculate cells with axons less than a millimeter long do not leave the geniculate but synapse locally on other geniculate cells. Despite these complicating features, single cells in the geniculate respond to light in much the same way as retinal ganglion cells, with similar on-center and off-center receptive fields and similar responses to color. In terms of visual information, then, the lateral geniculate bodies do not seem to be exerting any profound transformation, and we simply don't yet know what to make of the nonvisual inputs and the local synaptic interconnection.

LEFT AND RIGHT IN THE VISUAL PATHWAY

The optic fibers distribute themselves to the two lateral geniculate bodies in a special and, at first glance, strange way. Fibers from the left half of the left retina go to the geniculate on the same side, whereas fibers from the left half of the right retina cross at the optic chiasm and go to the opposite geniculate, as shown in the figure on page 2; similarly, the output of the two right half-retinas ends up in the right hemisphere. Because the retinal images are reversed by the lenses, light coming from anywhere in the right half of the visual environment projects onto the two left half-retinas, and the information is sent to the left hemisphere.

The term *visual fields* refers to the outer world, or visual environment, as seen by the two eyes. The *right visual field* means all points to the right of a vertical line through any point we are looking at, as illustrated in the diagram on this page. It is important to distinguish between *visual fields*, or what we see in the external world, and *receptive field*, which means the outer world as seen by a single cell. To reword the previous paragraph: the information from the right visual field projects onto the left hemisphere.




The right visual field extends out to the right almost to 90 degrees, as you can easily verify by wiggling a finger and slowly moving it around to your right. It extends up 60 degrees or so, down perhaps 75 degrees and to the left, by definition, to a vertical line passing through the point you are looking at.

Much of the rest of the brain is arranged in an analogous way: for example, information about touch and pain coming from the right half of the body goes to the left hemisphere; motor control to the right side of the body comes from the left hemisphere. A massive stroke in the left side of the brain leads to paralysis and lack of sensation in the right face, arm, and leg and to loss of speech. What is less commonly known is that such a stroke generally leads also to blindness in the right half of the visual world—the right visual field— involving *both* eyes. To test for such blindness, the neurologist has the patient stand in front of him, close one eye, and look at his (the neurologist's) nose with the other eye. He then explores the patient's visual fields by waving his hand or holding a Q-tip here and there and, in the case of a left-sided stroke, can show that the patient does not see anything to the right of where he is looking. For example, as the Q-tip is held up in the air between patient and neurologist, a bit above their heads, and moved slowly from the patient's right to left, the patient sees nothing until the white cotton crosses the midline, when it suddenly appears. The result is exactly the same when the other eye is tested. A complete right homonymous hemianopia (as neurologists call this are looking at half-blindness!) actually dissects precisely the foveal region (the center of gaze): if you look at the word *was*, riveting your gaze on the *a*, you won't see the *s*, and you will only see the left half of the *a*—an interesting if distressing


experience. We can see from such tests that each eye sends information to both hemispheres or, conversely, that each hemisphere of the brain gets input from both eyes. That may seem surprising. After my remarks about touch and pain sensation and motor control, you might have expected that the left eye would project to the right hemisphere and vice versa. But each hemisphere of the brain is dealing with the opposite half of the *environment*, not with the opposite side of the body. In fact, for the left eye to project to the right hemisphere is roughly what happens in many lower mammals such as horses and mice, and exactly what happens in lower vertebrates such as birds and amphibia. In horses and mice the eyes tend to point outward rather than straight ahead, so that most of the retina of the right eye gets its information from the right visual field, rather than from both the left and right visual fields, as is the case in forward-looking primates like ourselves. The description I have given of the visual pathway applies to mammals such as primates, whose two eyes point more or less straight ahead and therefore take in almost the same scene. Hearing works in a loosely analogous way. Obviously, each ear is capable of taking in sound coming from either the left or the right side of a person's world. Like each eye, each ear sends auditory information to both halves of the brain roughly equally, but in hearing, as in vision, the process still tends to be lateralized: sound coming to someone's two ears from any point to the right of where he or she is facing is processed, in the brainstem, by comparing amplitudes and time of arrival at the two ears, in such a way that the responses are for the most part channeled to higher centers on the left side. I am speaking here of early stages of information handling. By speaking or gesturing, someone standing to my right can persuade me to move my left hand, so that the information sooner or later has to get to my right hemisphere, but it must come first to my left auditory or visual cortex. Only then does it cross to the motor cortex on the other side. Incidentally, no one knows why the right half of the world tends to project to the left half of the cerebral hemispheres. The rule has important exceptions: the hemispheres of our cerebellum (a part of the brain concerned largely with movement control) get input largely from the same, not the opposite, half of the world. That complicates things for the brain, since the fibers connecting the cerebellum on one side to the motor part of the cerebral cortex on the other all have to cross from one side to the other. All that can be said with assurance is that this pattern is mysterious.

LAYERING OF THE LATERAL GENICULATE

Each lateral geniculate body is composed of six layers of cells stacked one on the other like a club sandwich. Each layer is made up of cells piled four to ten or more deep. The whole sandwich is folded along a fore-and-aft axis, giving the cross-sectional appearance shown in the illustration on the top of the next page.



The six cell layers show clearly in the left lateral geniculate body of a macaque monkey, seen in a section cut parallel to the face. The section is stained to show cell bodies, each of which appears as a dot.




The stacked-plate organization is preserved in going from retina to geniculate, except that the fibers from the retinas are bundled into a cable and splayed out again, in an orderly way, at their geniculate destination.

In the scheme in which one plate projects to the next, an important complication arises in the transition from retina to geniculate; here the two eyes join up, with the two separate plates of retinal ganglion cells projecting to the sextuple geniculate plate. A single cell in the lateral

geniculate body does not receive convergent input from the two eyes: a cell is a right-eye cell or a left-eye cell. These two sets of cells are segregated into separate layers, so that all the cells in any one layer get input from one eye only. The layers are stacked in such a way that the eyes alternate. In the left lateral geniculate body, the sequence in going from layer to layer, from above downwards, is right, left, right, left, left, right. It is not at all clear why the sequence reverses between the fourth and fifth layers—sometimes I think it is just to make it harder to remember. We really have no good idea why there is a sequence at all. As a whole, the sextuple-plate structure has just one topography. Thus the two left half-retinal surfaces project to one sextuple plate, the left lateral geniculate (see the bottom figure on the previous page). Similarly, the right half-retinas project to the right geniculate. Any single point in one layer corresponds to a point in the layer implies movement in the visual field along some path dictated by the visual-field-to-geniculate map. If we move instead in a direction perpendicular to the layers—for example, along the radial line in the figure on the top of the previous page—as the electrode passes from one layer to the next, the receptive fields stay in the same part of the visual field but the eyes switch—except, of course, where the sequence reverses. The half visual field maps onto each geniculate six times, three for each eye, with the maps in precise register. The lateral geniculate body seems to be two organs in one. With some justification we can consider the ventral, or bottom, two layers (*ventral* means "belly") as an entity because the cells they contain are different from the cells in the other four layers: they are bigger and respond differently to visual stimuli. We should also consider the four *dorsal*, or upper, layers (*dorsal* means "back" as opposed to "belly") as a separate structure because they are histologically and physiologically so similar to each other. Because of the different sizes of their cells, these two sets of layers are called *magnocellular* (ventral) and *parvocellular* (dorsal). Fibers from the six layers combine in a broad band called the *optic radiations*, which ascends to the primary visual cortex (see the illustration on page 2.) There, the fibers fan out in a regular way and distribute themselves so as to make a single orderly map, just as the optic nerve did on reaching the geniculate. This brings us, finally, to the cortex.

RESPONSES OF CELLS IN THE CORTEX


The main subject of this chapter is how the cells in the primary visual cortex respond to visual stimuli. The receptive fields of lateral geniculate cells have the same center-surround organization as the retinal ganglion cells that feed into them. Like retinal ganglion cells, they are distinguishable from one another chiefly by whether they have on centers or off centers, by their positions in the visual field, and by their detailed color properties. The question we now ask is whether cortical cells have the same properties as the geniculate cells that feed them, or whether they do something new. The answer, as you must already suspect, is that they indeed do something new, something so original that prior to 1958, when cortical cells were first studied with patterned light stimulation, no one had remotely predicted it.



This Golgi-stained section of the primary visual cortex shows over a dozen pyramidal cells—still just a tiny fraction of the total number in such a section. The height of the section is about 1 millimeter. (The long trunk near the right edge is a blood vessel.)

The primary visual, or striate, cortex is a plate of cells 2 millimeters thick, with a surface area of a few square inches. Numbers may help to convey an impression of the vastness of this structure: compared with the geniculate, which has 1.5 million cells, the striate cortex contains something like 200 million cells. Its structure is intricate and fascinating, but we don't need to know the details to appreciate how this part of the brain transforms the incoming visual information. We will look at the anatomy more closely when I discuss functional architecture in the next chapter. I have already mentioned that the flow of information in the cortex takes place over several loosely defined stages. At the first stage, most cells respond like geniculate cells. Their receptive fields have circular symmetry, which means that a line or edge produces the same response regardless of how it is oriented. The tiny, closely packed cells at this stage are not easy to record from, and it is still unclear whether their responses differ at all from the responses of geniculate cells, just as it is unclear whether the responses of retinal ganglion cells and geniculate cells differ. The complexity of the histology (the microscopic anatomy) of both geniculate and cortex certainly leads you to expect differences if you compare the right things, but it can be hard to know just what the "right things" are. This point is even more important when it comes to the responses of the cells at the next stage in the cortex, which presumably get their input from the center-surround cortical cells in the first stage. At first, it was not at all easy to know what these second-stage cells responded to. By the late 1950s very few scientists had attempted to record from single cells in the visual cortex, and those who did had come up with disappointing results. They found that cells in the visual cortex seemed to work very much like cells in the retina: they found on cells and off cells, plus an additional class that did not seem to respond to light at all. In the face of the obviously fiendish complexity of the cortex's anatomy, it was puzzling to find the physiology so boring. The explanation, in retrospect, is very clear. First, the stimulus was inadequate: to activate cells in the cortex, the usual custom was simply to flood the retina with diffuse light, a stimulus that is far from optimal even in the retina, as Kuffler had shown ten years previously. For most cortical cells, flooding the retina in this way is not only not optimal—it is completely without effect. Whereas many geniculate cells respond to diffuse white light, even if weakly, cortical cells, even those first-stage cells that resemble geniculate cells, give virtually no responses. One's first intuition, that the best way to activate a visual cell is to activate all the receptors in the retina, was evidently seriously off the mark. Second, and still more ironic, it turned out that the cortical cells that did give on or off responses were in fact not cells at all but merely axons coming in from the lateral geniculate body. The cortical cells were not responding at all! They were much too choosy to pay attention to anything as crude as diffuse light. This was the situation in 1958, when Torsten Wiesel and I made one of our first technically successful recordings from the cortex of a cat. The

position of microelectrode tip, relative to the cortex, was unusually stable, so much so that we were able to listen in on one cell for a period of about nine hours. We tried everything short of standing on our heads to get it to fire. (It did fire spontaneously from time to time, as most cortical cells do, but we had a hard time convincing ourselves that our stimuli had caused any of that activity.) After some hours we began to have a vague feeling that shining light in one particular part of the retina was evoking some response, so we tried concentrating our efforts there. To stimulate, we were using mostly white circular spots and black spots. For black spots, we would take a 1-by-2-inch glass microscope slide, onto which we had glued an opaque black dot, and shove it into a slot in the optical instrument Samuel Talbot had designed to project images on the retina. For white spots, we used a slide of the same size made of brass with a small hole drilled through it. (Research was cheaper in those days.) After about five hours of struggle, we suddenly had the impression that the glass with the dot was occasionally producing a response, but the response seemed to have little to do with the dot. Eventually we caught on: it was the sharp but faint shadow cast by the edge of the glass as we slid it into the slot that was doing the trick. We soon convinced ourselves that the edge worked only when its shadow was swept across one small part of the retina and that the sweeping had to be done with the edge in one particular orientation. Most amazing was the contrast between the machine-gun discharge when the orientation of the stimulus was just right and the utter lack of a response if we changed the orientation or simply shined a bright flashlight into the cat's eyes.



Responses of one of the first orientation-specific cells. Torsten Wiesel and I recorded, from a cat striate cortex in 1958. This cell not only responds exclusively to a moving slit in an eleven o'clock orientation but also responds to movement right and up, but hardly at all to movement left and down.

The discovery was just the beginning, and for some time we were very confused because, as luck would have it, the cell was of a type that we came later to call complex, and it lay two stages beyond the initial, center-surround cortical stage. Although complex cells are the commonest type in the striate cortex, they are hard to comprehend if you haven't seen the intervening type. Beyond the first, center-surround stage, cells in the monkey cortex indeed respond in a radically different way. Small spots generally produce weak responses or none. To evoke a response, we first have to find the appropriate part of the visual field to stimulate, that is, the appropriate part of the screen that the animal is facing: we have to find the receptive field of the cell. It then turns out that the most effective way to influence a cell is to sweep some kind of line across the receptive field, in a

direction perpendicular to the line's orientation. The line can be light on a dark background (a slit) or a dark bar on a white background or an edge boundary between dark and light. Some cells prefer one of these stimuli over the other two, often very strongly; others respond about equally well to all three types of stimuli. What is critical is the orientation of the line: a typical cell responds best to some optimum stimulus orientation; the response, measured in the number of impulses as the receptive field is crossed, falls off over about 10 to 20 degrees to either side of the optimum, and outside that range it declines steeply to zero (see the illustration on the previous page). A range of 10 to 20 degrees may seem imprecise, until you remember that the difference between one o'clock and two o'clock is 30 degrees. A typical orientation-selective cell does not respond at all when the line is oriented 90 degrees to the optimal. Unlike cells at earlier stages in the visual path, these orientation-specific cells respond far better to a moving than to a stationary line. That is why, in the diagram on the next page, we stimulate by sweeping the line over the receptive field. Flashing a stationary line on and off often evokes weak responses, and when it does, we find that the preferred orientation is always the same as when the line is moved. In many cells, perhaps one-fifth of the population, moving the stimulus brings out another kind of specific response. Instead of firing equally well to both movements, back and forth, many cells will consistently respond better to one of the two directions. One movement may even produce a strong response and the reverse movement none or almost none, as illustrated in the figure on the next page. In a single experiment we can test the responses of 200 to 300 cells simply by learning all about one cell and then pushing the electrode ahead to the next cell to study it. Because once you have inserted the delicate electrode you obviously can't move it sideways without destroying it or the even more delicate cortex, this technique limits your examination to cells lying in a straight line. Fifty cells per millimeter of penetration is about the maximum we can get with present methods. When the orientation preferences of a few hundred or a thousand cells are examined, all orientations turn out to be about equally represented—vertical, horizontal, and every possible oblique. Considering the nature of the world we look at, containing as it does trees and horizons, the question arises whether any particular orientations, such as vertical and horizontal, are better represented than the others. Answers differ with different laboratory results, but everyone agrees that if biases do exist, they must be small—small enough to require statistics to discern them, which may mean they are negligible! In the monkey striate cortex, about 70 to 80 percent of cells have this property of orientation specificity. In the cat, all cortical cells seem to be orientation selective, even those with direct geniculate input. We find striking differences among orientation-specific cells, not just in optimum stimulus orientation or in the position of the receptive field on the retina, but in the way cells behave. The most useful distinction is between two classes of cells: *simple* and *complex*. As their names suggest, the two types differ in the complexity of their behavior, and we make the reasonable assumption that the cells with the simpler behavior are closer in the circuit to the input of the cortex.

SIMPLE CELLS


For the most part, we can predict the responses of simple cells to complicated shapes from their responses to small-spot stimuli. Like retinal ganglion cells, geniculate cells, and circularly symmetric cortical cells, each simple cell has a small, clearly delineated receptive field within which a small spot of light produces either on or off responses, depending on where in the field the spot falls. The difference between these cells and cells at earlier levels is in the geometry of

the maps of excitation and inhibition. Cells at earlier stages have maps with circular symmetry, consisting of one region, on or off, surrounded by the opponent region, off or on. Cortical simple cells are more complicated.


The excitatory and inhibitory domains are always separated by a straight line or by two parallel lines, as shown in the three drawings on the next page. Of the various possibilities, the most common is the one in which a long, narrow region giving excitation is flanked on both sides by larger regions giving inhibition, as shown in the first drawing (a). To test the predictive value of the maps made with small spots, we can now try other shapes. We soon learn that the more of a region a stimulus fills, the stronger is the resultant excitation or inhibition; that is, we find *spatial* summation of effects. We also find *antagonism*, in which we get a mutual cancellation of responses on stimulating two opposing regions at the same time. Thus for a cell with a receptive-field map like that shown in the first drawing (a), a long, narrow slit is the most potent stimulus, provided it is positioned and oriented so as to cover the excitatory part of the field without invading the inhibitory part (see the illustration on the next other page.) Even the slightest part (see the illustration on the next page.) Even the slightest misorientation causes the slit to miss some of the excitatory area and to invade the antagonistic inhibitory part, with a consequent decline in response.

In the second and third figures (b and c) of the diagram on the next page, we see two other kinds of simple cells: these respond best to dark lines and to dark/ light edges, with the same sensitivity to the orientation of the stimulus. For all three types, diffuse light evokes no response at all. The mutual cancellation is obviously very precise, reminiscent of the acid-base titrations we all did in high school chemistry labs. Already, then, we can see a marked diversity in cortical cells. Among simple cells, we find three or four different geometries, for each of which we find every possible orientation and all possible visual-field positions.

The size of a simple-cell receptive field depends on its position in the retina relative to the fovea, but even in a given part of the retina, we find some variation in size. The smallest fields, in and near the fovea, are about one-quarter degree by one-quarter degree in total size; for a cell of the type shown in diagrams a or b in the figure on the next page, the center region has a width of as little as a few minutes of arc. This is the same as the diameters of the smallest receptive-field centers in retinal ganglion cells or geniculate cells. In the far retinal periphery, simple-cell receptive fields can be about 1 degree by 1 degree.



Three typical receptive-field maps for simple cells. The effective stimuli for these cells are (a) a slit covering the plus (+) re- gion, (b) a dark line covering the minus (—) region, and (c) a light-dark edge falling on the boundary between plus and minus.




Various stimulus geometries evoke different responses in a cell with receptive field of the type in diagram a of the previous figure. The stimulus line at the bottom indicates when the slit is turned on and, 1 second later, turned off. The top record shows the response to a slit of optimum size, position, and orientation. In the second record, the same slit covers only part of an inhibitory area. (Because this cell has no spontaneous activity to suppress, only an off discharge is seen.) In the third record, the slit is oriented so as to cover only a small part of the excitatory region and a proportionally small part of the inhibitory region; the cell fails to respond. In the bottom record, the whole receptive field is illuminated; again, there is no response.

Even after twenty years we still do not know how the inputs to cortical cells are wired in order to bring about this behavior. Several plausible circuits have been proposed, and it may well be that one of them, or several in combination, will turn out to be correct. Simple cells must be built up from the antecedent cells with circular fields; by far the simplest proposal is that a simple cell receives direct excitatory input from many cells at the previous stage, cells whose receptive-field centers are distributed along a line in the visual field, as shown in the diagram on the next page.

It seems slightly more difficult to wire up a cell that is selectively responsive to edges, as shown in the third drawing (c) on the previous page. One workable scheme would be to have the cell receive inputs from two sets of antecedent cells having their field centers arranged on opposite sides of a line, on-center cells on one side off-center cells on the other all making excitatory connections cells on one side, off-center cells on the other, all making excitatory connections. In all these proposed circuits, excitatory input from an off-center cell is logically equivalent to inhibitory input from an on-center cell, provided we assume that the off-center cell is spontaneously active.

Working out the exact mechanism for building up simple cells will not be easy. For any one cell we need to know what kinds of cells feed in information—for example, the details of their receptive fields, including position, orientation if any, and whether on or off center—and whether they supply excitation or inhibition to the cell. Because methods of obtaining this kind of knowledge don't yet exist, we are forced to use less direct approaches, with correspondingly higher chances of being wrong. The mechanism summarized in the diagram on the next page seems to me the most likely because it is the most simple.





This type of wiring could produce a simple-cell receptive field. On the right, four cells are shown making excitatory synaptic connections with a cell of higher order. Each of the lower-order cells has a radially symmetric receptive field with on- center and off-surround, illustrated by the left side of the diagram. The centers of these fields lie along a line. If we suppose that many more than four center-surround cells are connected with the simple cell, all with their field centers overlapped along this line, the receptive field of the simple cell will consist of a long, narrow excitatory region with inhibitory flanks. Avoiding receptive-field terminology, we can say that stimulating with a small spot anywhere in this long, narrow rectangle will strongly activate one or a few of the center-surround cells and in turn excite the simple cell, although only weakly. Stimulating with a long, narrow slit will activate all the center-surround cells, producing a strong response in the simple cell.

COMPLEX

CELLS


Complex cells represent the next step or steps in the analysis. They are the commonest cells in the striate cortex—a guess would be that they make up three-quarters of the population. The first oriented cell Wiesel and I recorded—the one that responded to the edge of the glass slide—was in retrospect almost certainly a complex cell.

Complex cells share with simple cells the quality of responding only to specifically oriented lines. Like simple cells, they respond over a limited region of the visual field; unlike simple cells, their behavior cannot be explained by a neat subdivision of the receptive field into excitatory and inhibitory regions. Turning a small stationary spot on or off seldom produces a response, and even an appropriately oriented stationary slit or edge tends to give no response or only weak, unsustained responses of the same type everywhere—at the onset or turning off of the stimulus or both. But if the properly oriented line is swept across the receptive field, the result is a well-sustained barrage of impulses, from the instant the line enters the field until it leaves (see the cell-response diagram on page 10). By contrast, to evoke sustained responses from a simple cell, a stationary line must be critically oriented and critically positioned; a moving line evokes only a brief response at the moment it crosses a boundary from an inhibitory to an excitatory region or during the brief time it covers the excitatory region. Complex cells that do react to stationary slits, bars, or edges fire regardless of where the line is placed in the receptive field, as long as the orientation is appropriate. But over the same region, an inappropriately oriented line is ineffective, as shown in the illustration on the next page.



Left: A long, narrow slit of light evokes a response wherever it is placed within the receptive field (rectangle) of a complex cell, provided the orientation is correct (upper three records). A nonoptimal orientation gives a weaker response or none at all (lower record). *Right*: The cortical cell from layer 5 in the striate cortex of a cat was recorded intracellularly by David Van Essen and James Kelly at Harvard Medical School in 1973, and its complex receptive field was mapped. They then injected procyon yellow dye and showed that the cell was pyramidal.

The diagram on this page for the complex cell and the one on page 13 for the simple cell illustrate the essential difference between the two types: for a simple cell, the extremely narrow range of positions over which an optimally oriented line evokes a response; for a complex cell, the responses to a properly oriented line wherever it is placed in the receptive field. This behavior is related to the explicit on and off regions of a simple cell and to the lack of such regions in a complex cell. The complex cell generalizes the responsiveness to a line over a wider territory. Complex cells tend to have larger receptive fields than simple cells, but not very much larger. A typical complex receptive field in the fovea of the macaque monkey would be about one-half degree by one-half degree. The optimum stimulus width is about the same for simple cells and complex cells—in the fovea, about 2 minutes of arc. The complex cell's resolving power, or acuity, is thus the same as the simple cell's. As in the case of the simple cell, we do not know exactly how complex cells are built up. But, again, it is easy to propose plausible schemes, and the simplest one is to imagine that the complex cell receives input from many simple cells, all of whose fields have the same orientation but are spread out in overlapping fashion over the entire field of the complex cell, as shown in the illustration on the next page. If the connections from simple to complex cells are excitatory, then wherever a line falls in the field, *some* simple cells are activated; the complex cell will therefore be activated. If, on the other hand, a stimulus fills the entire receptive field, none of the simple cells will be activated, and the complex cell won't be activated.




This wiring diagram would account for the properties of a complex cell. As in the figure on page 15, we suppose that a large number of simple cells (only three are shown here) make excitatory synapses with a single complex cell. Each simple cell responds optimally to a vertically oriented edge with light to the right, and the receptive fields are scattered in overlapping fashion throughout the rectangle. An edge falling anywhere within the rectangle evokes a response from a few simple cells, and this in turn evokes a response in the complex cell. Because there is adaptation at the synapses, only a moving stimulus will keep up a steady bombardment of the complex cell.

The burst of impulses from a complex cell to turning on a stationary line and not moving it is generally brief even if the light is kept on: we say that the response *adapts*. When we move the line through the complex cell's receptive field, the sustained response may be the result of overcoming the adaptation, by bringing in new simple cells one after the next. You will have noticed that the schemes for building simple cells from center-surround ones, as in the illustration on page 15, and for building complex cells out of simple ones, as in the illustration on this page, both involve excitatory processes. In the two cases, however, the processes must be very different. The first scheme requires *simultaneous* summed inputs from centersurround cells whose field centers lie along a line. In the second scheme, activation of the complex cell by a moving stimulus requires *successive* activation of many simple cells. It would be interesting to know what, if any, morphological differences underlie this difference in addition properties.

DIRECTIONAL SELECTIVITY

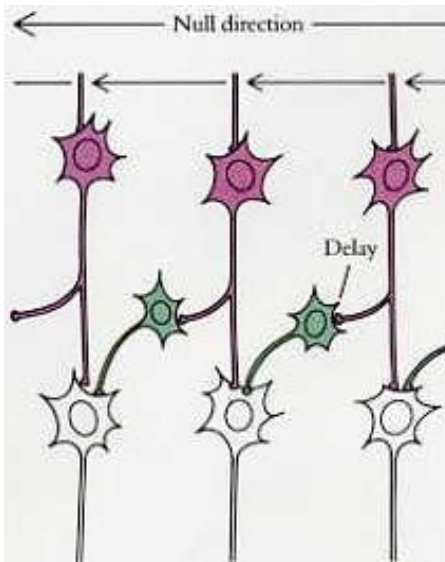
Many complex cells respond better to one direction of movement than to the diametrically opposite direction. The difference in response is often so marked that one direction of movement will produce a lively response and the other direction no response at all, as shown in the diagram on the next page. It turns out that about 10 to 20 percent of cells in the upper layers of the striate cortex show marked directional selectivity. The rest seem not to care: we have to pay close attention or use a computer to see any difference in the responses to the two opposite directions. There seem to be two distinct classes of cells, one strongly direction-selective, the other not selective. Listening to a strongly direction-selective cell respond, the feeling you get is that the line moving in one direction grabs the cell and pulls it along and that the line moving in the other direction fails utterly to engage it—something like the feeling you get with a ratchet, in winding a watch.



Responses of this complex cell differ to an optimally oriented slit moving in opposite directions. Each record is about 2 seconds in duration. (Cells such as this are not very fussy about how fast the slit moves; generally, responses fail only when the slit moves so fast that it becomes blurred or so slow that movement cannot be seen.)

We don't know how such directionally selective cells are wired up. One possibility is that they are built up from simple cells whose responses to opposite directions of movement are asymmetric. Such simple cells have asymmetric fields, such as the one shown in the third diagram of the illustration on page 13. A second mechanism was proposed in 1965 by Horace Barlow and William Levick to explain the directional selectivity of certain cells in the rabbit retina cells that apparently are not present in the monkey retina. If we apply their scheme to complex cells, we would suppose that interposed between simple and complex cells are intermediate cells, colored white in the diagram on the next page. We imagine that an intermediate cell receives excitation from one simple cell and inhibition from another (green) cell, whose receptive field is immediately adjacent and always located to one side and not the other. We further suppose that the inhibitory path involves a delay, perhaps produced by still another intermediate cell. Then, if the stimulus moves in one direction, say, right to left, as in the illustration of Barlow and Levick's model, the intermediate cell is excited by one of its inputs just as the inhibition arrives from the other, whose field has just been crossed. The two effects cancel, and the cell does not fire. For left-to-right movement, the inhibition arrives too late to prevent firing. If many such intermediate cells converge on a third cell, that cell will have the properties of a directionally selective complex cell. We have little direct evidence for any schemes that try to explain the behavior of cells in terms of a hierarchy of complexity, in which cells at each successive level are constructed of building blocks from the previous level. Nevertheless, we have strong reasons for believing that the nervous system is organized in a hierarchical series. The strongest evidence is anatomical: for example, in the cat, simple cells are aggregated in the fourth layer of the striate cortex, the same layer that receives geniculate input, whereas the complex cells are located in the layers above and below, one or two synapses further along. Thus although we may not know the exact circuit diagram at each stage, we have good reasons to suppose the existence of some circuit. The main reason for thinking that complex cells may be built up from center-surround cells, with a step in between, is the seeming necessity of doing the job in two logical steps. I should emphasize the word logical because the whole transformation presumably could be accomplished in one physical step by having center-surround inputs sum on separate dendritic branches of complex cells, with each branch doing the job of a simple cell, signaling electrotonically (by

passive electrical spread) to the cell body, and hence to the axon, whenever a line falls in some particular part of the receptive field. The cell itself would then be complex. But the very existence of simple cells suggests that we do not have to imagine anything as complicated as this.



Horace Barlow and William Levick proposed this circuit to explain directional selectivity. Synapses from purple to green are excitatory, and from green to white, inhibitory. We suppose the three white cells at the bottom converge on a single master cell.



THE SIGNIFICANCE OF MOVEMENT-SENSITIVE CELLS, INCLUDING SOME COMMENTS ON HOW WE SEE

Why are movement-sensitive cells so common? An obvious first guess is that they tell us if the visual landscape contains a moving object. To animals, ourselves included, changes in the outside world are far more important than static conditions, for the survival of predator and prey alike. It is therefore no wonder that most cortical cells respond better to a moving object than to a stationary one. Having carried the logic this far, you may now begin to wonder how we analyze a stationary landscape at all if, in the interests of having high movement sensitivity, so many oriented cells are insensitive to stationary contours. The answer requires a short digression, which takes us to some basic, seemingly counterintuitive facts about how we see. First, you might expect that in exploring our visual surroundings, we let our eyes freely rove around in smooth, continuous movement. What our two eyes in fact do is fixate on an object: we first adjust the positions of our eyes so that the images of the object fall on the two foveas; then we hold that position for a brief period, say, half a second; then our eyes suddenly jump to a new position by fixating on a new target whose presence somewhere out in the visual field has asserted itself, either by moving slightly, by contrasting with the background, or by presenting an interesting shape. During the jump, or saccade, which is French for "jolt", or "jerk" (the verb), the eyes move so rapidly that our visual system does not even respond to the resulting movement of the scene across the retina; we are altogether unaware of the violent change. (Vision may also in some sense be turned off during saccades by a complex circuit linking eye-movement centers with the visual

path.) Exploring a visual scene, in reading or just looking around, is thus a process of having our eyes jump in rapid succession from one place to another. Detailed monitoring of eye movements shows vividly how unaware we are of any of this. To monitor eye movements we first attach a tiny mirror to a contact lens, at the side, where it does not block vision; we then reflect a spot of light off the mirror onto a screen. Or, using a more modern version invented by David Robinson at the Wilmer Institute at Johns Hopkins, we can mount a tiny coil of wire around the rim of a contact lens, with the subject seated between two orthogonal pairs of bicycle-wheel size hoops containing coils of wire; currents in these coils induce currents in the contact-lens coil, which can be calibrated to give precise monitoring of eye movements. Neither method is what you would call a picnic for the poor subject. In 1957, Russian psychophysicist A. L. Yarbus recorded eye movements of subjects as they explored various images, such as a woods or female faces (see the illustrations below), by showing the stopping places of a subject's gaze as dots joined by lines indicating the eyes' trajectory during the jumps. A glance at these amazing pictures gives us a world of information about our vision—even about the objects and details that interest us in our environment. So the first counterintuitive fact is that in visual exploration our eyes jump around from one point of interest to another: we cannot explore a stationary scene by swinging our eyes



past it in continuous movements.



A picture is viewed by an observer while we monitor eye position and hence direction of gaze. The eyes jump, come to rest momentarily (producing a small dot on the record), then jump to a new locus of interest. It seems difficult to jump to a void—a place lacking abrupt luminance changes.

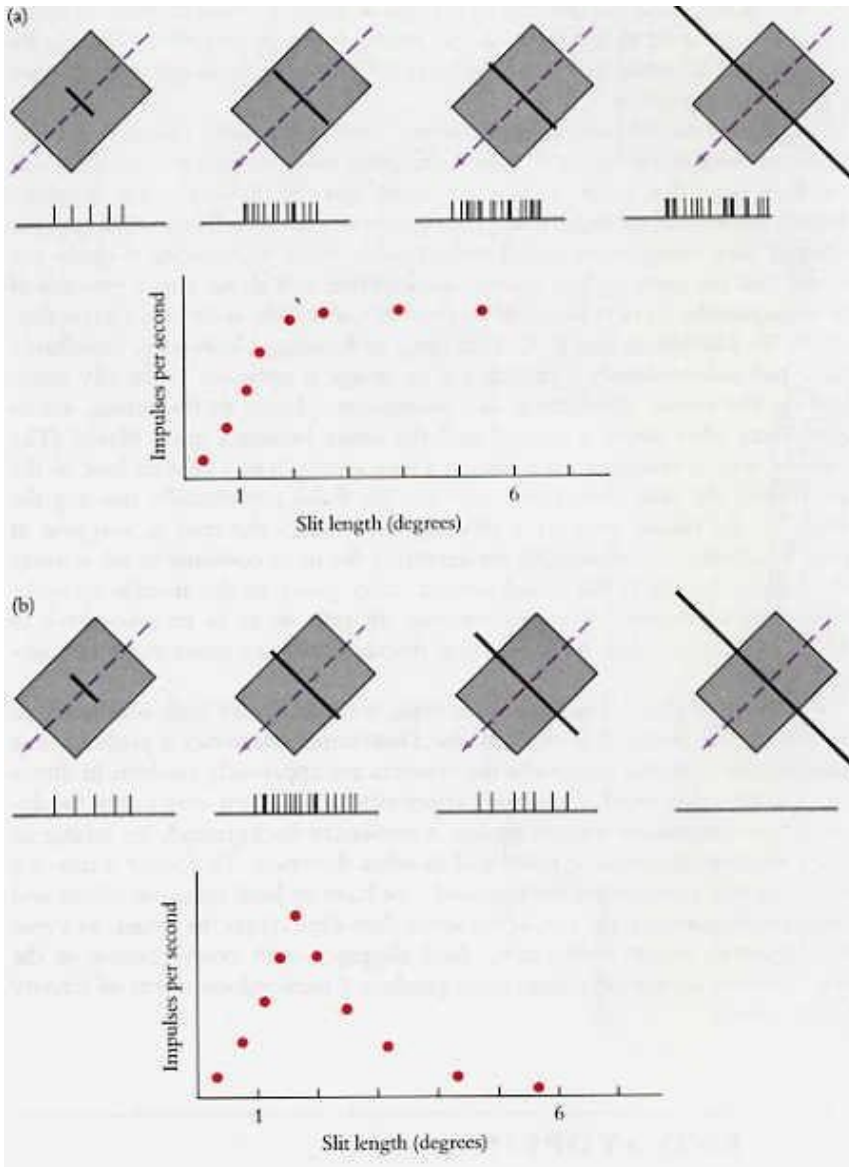
The visual system seems intent instead on keeping the image of a scene anchored on our retinas, on preventing it from sliding around. If the whole scene moves by, as occurs when we look out a train window, we follow it by fixating on an object and maintaining fixation by moving our eyes until the object gets out of range, whereupon we make a saccade to a new object. This whole sequence— following with smooth pursuit, say, to the right, then making a saccade to the left—is called *nystagmus*. You can observe the sequence—

perhaps next time you are in a moving train or streetcar—by looking at your neighbor's eyes as he or she looks out a window at the passing scene—taking care not to have your attentions misunderstood! The process of making visual saccades to items of interest, in order to get their images on the fovea, is carried out largely by the superior colliculus, as Peter Schiller at MIT showed in an impressive series of papers in the 1970s. The second set of facts about how we see is even more counterintuitive. When we look at a stationary scene by fixating on some point of interest, our eyes lock onto that point, as just described, but the locking is not absolute. Despite any efforts we may make, the eyes do not hold perfectly still but make constant tiny movements called *microsaccades*; these occur several times per second and are more or less random in direction and about 1 to 2 minutes of arc in amplitude. In 1952 Lorrin Riggs and Floyd Ratliff, at Brown University, and R. W. Ditchburn and B. L. Ginsborg, at Reading University, simultaneously and independently found that if an image is optically artificially stabilized on the retina, eliminating any movement relative to the retina, vision fades away after about a second and the scene becomes quite blank! (The simplest way of stabilizing is to attach a tiny spotlight to a contact lens; as the eye moves, the spot moves too, and quickly fades.) Artificially moving the image on the retina, even by a tiny amount, causes the spot to reappear at once. Evidently, microsaccades are necessary for us to continue to see stationary objects.

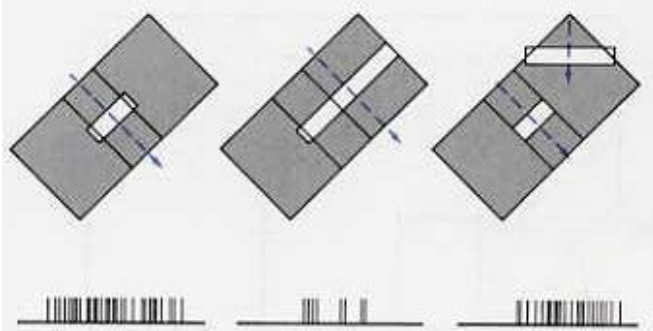
It is as if the visual system, after going to the trouble to make movement a powerful stimulus—wiring up cells so as to be insensitive to stationary objects—had then to invent microsaccades to make stationary objects visible. We can guess that cortical complex cells, with their very high sensitivity to movement, are involved in this process. Directional selectivity is probably not involved, because microsaccadic movements are apparently random in direction. On the other hand, directional selectivity would seem very useful for detecting movements of objects against a stationary background, by telling us that a movement is taking place and in what direction. To follow a moving object against a stationary background, we have to lock onto the object and track it with our eyes; the rest of the scene then slips across the retina, an event that otherwise occurs only rarely. Such slippage, with every contour in the scene moving across the retina, must produce a tremendous storm of activity in our cortex.

END STOPPING

One additional kind of specificity occurs prominently in the striate cortex. An ordinary simple or complex cell usually shows length summation: the longer the stimulus line, the better is the response, until the line is as long as the receptive field; making the line still longer has no effect. For an *end-stopped cell*, lengthening the line improves the response up to some limit, but exceeding that limit in one or both directions results in a weaker response, as shown in the bottom diagram on the next page.




Top: An ordinary complex cell responds to various lengths of a slit of light. The duration of each record is 2 seconds. As indicated by the graph of response versus slit length, for this cell the response increases with length up to about 2 degrees, after which there is no change. *Bottom:* For this end-stopped cell, responses improve up to 2 degrees but then decline, so that a line 6 degrees or longer gives no response.



For this end-stopped cell, stimulating the middle activating region alone with an optimally oriented slit produces a strong response. Including one of the inhibitory regions almost nullifies the response, but if the inhibitory region is stimulated with a different orientation, the response is no longer blocked. Thus the activating region and the inhibitory regions both have the same optimal orientations.

Some cells, which we call completely end stopped, do not respond at all to a long line. We call the region from which responses are evoked the activating region and speak of the regions at one or both ends as inhibitory. The total receptive field is consequently made up of the activating region and the inhibitory region or regions at the ends. The stimulus orientation that best evokes excitation in the activating region evokes maximal inhibition in the outlying area(s). This can be shown by repeatedly stimulating the *activating region* with an optimally oriented line of optimal length while testing the outlying region with lines of varying orientation, as shown in the top diagram on this page.




One scheme for explaining the behavior of a complex end-stopped cell. Three ordinary complex cells converge on the end-stopped cell: one, whose receptive field is congruent with the end-stopped cell's activating region (a), makes excitatory contacts; the other two, having fields in the outlying regions (b and c), make inhibitory contacts.

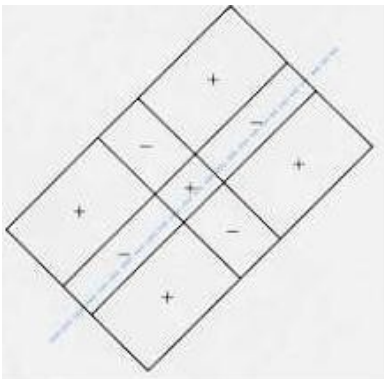
We originally thought that such cells represented a stage one step beyond complex cells in the hierarchy. In the simplest scheme for elaborating such a cell, the cell would be excited by one or a few ordinary complex cells with fields in the activating region and would be inhibited by complex cells with similarly oriented fields situated in the outlying regions. I have illustrated this scheme in the bottom diagram on this page. A second possibility is that the cell receives excitatory input from cells with small fields, marked (a) in the diagram on this page, and inhibition from cells with large fields, marked (b); we assume that the cells supplying inhibition are maximally excited by long slits but poorly excited by short ones. This second possibility (analogous to the model for center-

surround cells described in Section "Bipolar cells and horizontal cells" of Chapter 3) is one of the few circuits for which we have some evidence. Charles Gilbert, at Rockefeller University in New York, has shown that complex cells in layer 6 of the monkey striate cortex have just the right properties for supplying this inhibition and, furthermore, that disabling these cells by local injections causes end-stopped cells in the upper layers to lose the end inhibition.

After these models were originally proposed, Geoffery Henry, in Canberra, Australia, discovered end-stopped simple cells, presumably with the receptive-field arrangement shown in the lower margin of this page. For such a cell, the wiring would be analogous to our first diagram, except that the input would be from simple rather than from complex cells.




In an alternative scheme, one cell does the inhibiting, a cell whose receptive field covers the entire area (b). For this to work, we have to assume that the inhibiting cell responds only weakly to a short slit when (a) is stimulated, but responds strongly to a long slit.



This end-stopped simple cell is assumed to result from convergent input from three ordinary simple cells. (One cell, with the middle on-center field, could excite the cell in question; the two others could be off center and also excite or be on center and inhibit.) Alternatively, the input to this cell could come directly from center-surround cells, by some more elaborate version of the process illustrated on page 15.


Complex end-stopped cells could thus arise by excitatory input from one set of complex cells and inhibitory input from another set, as in the diagrams on the two preceding pages, or by convergent input from many end-stopped simple cells. The optimal stimulus for an end-stopped cell is a line that extends for a certain distance and no further. For a cell that responds to edges and is end stopped at one end only, a corner is ideal; for a cell that responds to slits or black bars and is stopped at both ends, the optimum stimulus is a short white or black line or a line that curves so that it is appropriate in the activating region and inappropriate—different by 20 to 30 degrees or more—in the flanking regions, as shown in the diagram of the curved contour on this page. We can thus view end-stopped cells as sensitive to corners, to curvature, or to sudden breaks in lines.



For an end-stopped cell such as the one shown on page 23, a curved border should be an effective stimulus.

THE IMPLICATIONS OF SINGLE-CELL PHYSIOLOGY FOR PERCEPTION

The fact that a cell in the brain responds to visual stimuli does not guarantee that it plays a direct part in perception. For example, many structures in the brainstem that are primarily visual have to do only with eye movements, pupillary constriction, or focusing by means of the lens. We can nevertheless be reasonably sure that the cells I described in this chapter have a lot to do with perception. As I mentioned at the outset, destroying any small piece of our striate cortex produces blindness in some small part of our visual world, and damaging the striate cortex has the same result in the monkey. In the cat things are not so simple: a cat with its striate cortex removed can see, though less well. Other parts of the brain, such as the superior colliculus, may play a relatively more important part in a cat's perception than they do in the primate's. Lower vertebrates, such as frogs and turtles, have nothing quite like our cortex, yet no one would contend that they are blind. We can now say with some confidence what any one of these cortical cells is likely to be doing in response to a natural scene. The majority of cortical cells respond badly to diffuse light but well to appropriately oriented lines. Thus for the kidney shape shown in the illustration on the next page, such a cell will fire if and only if its receptive field is cut in the right orientation by the borders. Cells whose receptive fields are inside the borders will be unaffected; they will continue to fire at their spontaneous rate, oblivious to the presence or absence of the form.



How arc cells in our brain likely to respond to some everyday stimulus, such as this kidney-shaped uniform blob? In the visual cortex, only a select set of cells will show any interest.

This is the case for orientation-specific cells in general. But to evoke a response from a simple cell, a contour must do more than be oriented to match the optimum orientation of the cell; it must also fall in the simple cell's receptive field, almost exactly on a border between excitation and inhibition, because the excitatory part must be illuminated without encroachment on the inhibitory part. If we move the contour even slightly, without rotating it, it will no longer stimulate the cell; it will now activate an entirely new population of simple cells. For complex cells, conditions are much less stringent because whatever population of cells is activated by a stimulus at one instant will remain unchanged if the form is moved a small distance in any direction without rotation. To cause a marked change in the population of activated complex cells, a movement has to be large enough for the border to pass entirely out of the receptive fields of some complex cells and into the fields of others. Thus compared to the population of simple cells, the population of activated complex cells, as a whole, will not greatly change in response to small translational movements of an object.

Finally, for end-stopped cells, we similarly find an increased freedom in the exact placement of the stimulus, yet the population activated by any form will be far more select. For end-stopped cells, the contour's orientation must fit the cell's optimum orientation within the activating region but must differ enough just beyond the activating region so as not to annul the excitation. In short, the contour must be curved just enough to fit the cell's requirements, or it must terminate abruptly, as shown in the diagram of the curve (page 25). One result of these exacting requirements is to increase efficiency, in that an object in the visual field stimulates only a tiny fraction of the cells on whose receptive fields it falls. The increasing cell specialization underlying this efficiency is likely to continue as we go further and deeper into the central nervous system, beyond the striate cortex. Rods and cones are influenced by light as such. Ganglion cells, geniculate cells, and center-surround cortical cells compare a region with its surrounds and are therefore likely to be influenced by any contours that cut their receptive fields but will not be influenced by overall changes in light intensity. Orientation-specific cells care not

only about the presence of a contour but also about its orientation and even its rate of change of orientation—its curvature. When such cells are complex, they are also sensitive to movement. We can see from the discussion in the last section that movement sensitivity can have two interpretations: it could help draw attention to moving objects, or it could work in conjunction with microsaccades to keep complex cells firing in response to stationary objects. I suspect light-dark contours are the most important component of our perception, but they are surely not the only component. The coloring of objects certainly helps in defining their contours, although our recent work tends to emphasize the limitations of color in defining forms. The shading of objects, consisting of gradual light-dark transitions, as well as their textures, can give important clues concerning shape and depth. Although the cells we have been discussing could conceivably contribute to the perception of shading and texture, we would certainly not expect them to respond to either quality with enthusiasm. How our brain handles textures is still not clear. One guess is that complex cells do mediate shades and textures without the help of any other

specialized sets of cells. Such stimuli may not activate many cells very efficiently, but the spatial extension that is an essential attribute of shading or texture may make many cells respond, all in a moderate or weak way. Perhaps lukewarm responses from many cells are enough to transmit the information to higher levels.

Many people, including myself, still have trouble accepting the idea that the interior of a form (such as the kidney bean on the previous page) does not itself excite cells in our brain—that our awareness of the interior as black or white (or colored, as we will see in Chapter 8) depends only on cells sensitive to the borders. The intellectual argument is that the perception of an evenly lit interior depends on the activation of cells having fields at the borders and on the absence of activation of cells whose fields are within the borders, since such activation would indicate that the interior is not evenly lit. So our perception of the interior as black, white, gray, or green has nothing to do with cells whose fields are in the interior—hard as that may be to swallow. But if an engineer were designing a machine to encode such a form, I think this is exactly what he would do. What happens at the borders is the only information you need to know: the interior is boring. Who could imagine that the brain would not evolve in such a way as to handle the information with the least number of cells?

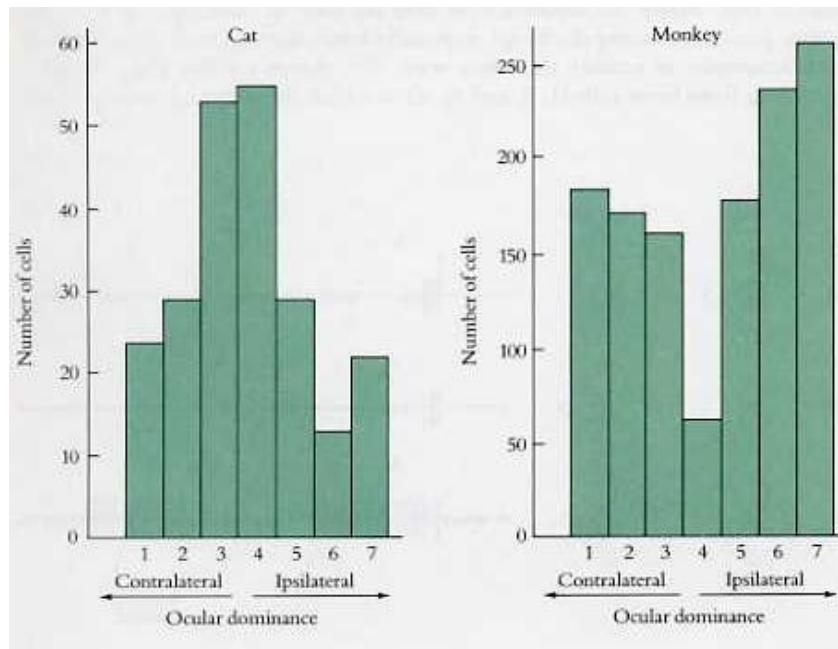
After hearing about simple and complex cells, people often complain that the analysis of every tiny part of our visual field—for all possible orientations and for dark lines, light lines, and edges—must surely require an astronomic number of cells. The answer is yes, certainly. But that fits perfectly, because an astronomic number of cells is just what the cortex has. Today we can say what the cells in this part of the brain are doing, at least in response to many simple, everyday visual stimuli. I suspect that no two striate cortical cells do exactly the same thing, because whenever a microelectrode tip succeeds in recording from two cells at a time, the two show slight differences—in exact receptive field position, directional selectivity, strength of response, or some other attribute. In short, there seems to be little redundancy in this part of the brain. How sure can we be that these cells are not wired up to respond to some other stimulus besides straight line segments? It is not as though we and others have not tried many other stimuli, including faces, *Cosmopolitan* covers, and waving our hands. Experience shows that we would be foolish to think that we had exhausted the list of possibilities. In the early 1960s, just when we felt satisfied with the striate cortex and wanted to move on to the next area (in fact, we had moved on), we happened to record from a sluggishly responding cell in the striate cortex and, by making the slit shorter, found that this very cell was anything but a sluggish responder. In this way we stumbled

on end stopping. And it took almost twenty years work with the monkey striate cortex before we became aware of blobs—pockets of cells specialized for color, described in Chapter 8. Having expressed these reservations, I should add that I have no doubt at all that some of the findings, such as orientation selectivity, are genuine properties of these cells. There is too much collateral evidence, such as the functional anatomy, described in Chapter 5, to allow for much scepticism.


BINOCULAR CONVERGENCE

I have so far made little mention of the existence of two eyes. Obviously, we need to ask whether any cortical cells receive input from both eyes and, if so, whether the two inputs are generally equal, qualitatively or quantitatively. To get at these questions we have to backtrack for a moment to the lateral geniculate body and ask if any of the cells there can be influenced from both eyes. The lateral geniculate body represents the first opportunity for information from the two eyes to come together at the level of the single cell. But it seems that the opportunity there is missed: the two sets of input are consigned to separate sets of layers, with little or no opportunity to combine. As we would expect from this segregation, a geniculate cell responds to one eye and not at all to the other. Some experiments have indicated that stimuli to the otherwise ineffective eye can subtly influence responses from the first eye. But for all practical purposes, each cell seems to be virtually monopolized by one or the other eye. Intuitively, it would seem that the paths from the two eyes must sooner or later converge, because when we look at a scene we see one unified picture. It is nevertheless everyone's experience that covering one eye makes no great difference in what we see: things seem about as clear, as vivid, and as bright. We see a bit farther to the side with both eyes, of course, because the retinas do not extend around as far in an outward (temporal) direction as they extend inwardly (nasally); still, the difference is only about 20 to 30 degrees. (Remember that the visual environment is inverted and reversed on the retina by the optics of the eye.) The big difference between one-eyed and two-eyed vision is in the sense of depth, a subject taken up in Chapter 7. In the monkey cortex, the cells that receive the input from the geniculatates, those whose fields have circular symmetry, are also like geniculate cells in being monocular. We find about an equal number of left-eye and right-eye cells, at least in parts of the cortex subserving vision up to about 20 degrees from the direction of gaze. Beyond this center-surround stage, however, we find binocular cells, simple and complex. In the macaque monkey over half of these higher-order cells can be influenced independently from the two eyes. Once we have found a binocular cell we can compare in detail the receptive fields in the two eyes. We first cover the right eye and map the cell's receptive field in the left eye, noting its exact position on the screen or retina and its complexity, orientation, and arrangement of excitatory and inhibitory regions; we ask if the cell is simple or complex, and we look for end stopping and directional selectivity. Now we block off the left eye and uncover the right, repeating all the questions. In most binocular cells, we find that all the properties found in the left eye hold also for the right-eye stimulation—the same position in the visual field, the same directional selectivity, and so on. So we can say that the connections or circuits between the left eye and the cell we are studying are present as a duplicate copy between the right eye and that cell. We need to make one qualification concerning this duplication of connections. If, having found the best stimulus—orientation, position, movement direction, and so on—we then compare the responses evoked

from one eye with the responses evoked from the other, we find that the two responses are not necessarily equally vigorous. Some cells do respond equally to the two eyes, but others consistently give a more powerful discharge to one eye than to the other. Overall, except for the part of the cortex subserving parts of the visual field well away from the direction of gaze, we find no obvious favoritism: in a given hemisphere, just as many cells favor the eye on the opposite side (the *contralateral* eye) as the eye on the same side (the *ipsilateral*). All shades of relative eye dominance are represented, from cells monopolized by the left eye through cells equally affected to cells responding only to the right eye. We can now do a population study. We group all the cells we have studied, say 1000 of them, into seven arbitrary groups, according to the relative effectiveness of the two eyes; we then compare their numbers, as shown in the two bar graphs on this page. At a glance the histograms tell us how the distribution differs between cat and monkey: that in both species, binocular cells are common, with each eye well represented (roughly equally, in the monkey); that in cats, binocular cells are very abundant; that in macaques, monocular and binocular cells are about equally common, but that binocular cells often favor one eye strongly (groups 2 and 5). We can go even further and ask if binocular cells respond better to both eyes than to one. Many do: separate eyes may do little or nothing, but both together produce a strong discharge, especially when the two eyes are stimulated simultaneously in exactly the same way. The figure on the next page shows a recording from three cells (1, 2, and 3), all of which show strong synergy.



In population studies of ocular dominance, we study hundreds of cells and categorize each one as belonging to one of seven arbitrary groups. A group *i* cell is defined as a cell influenced only by the contralateral eye—the eye opposite to the hemisphere in which it sits. A group 2 cell responds to both eyes but strongly prefers the contralateral eye. And so on.



The recording electrode was close enough to three cells to pick up impulses from all of them. Responses could be distinguished by size and shape of the impulses. This illustrates the responses to stimuli to single eyes and to both eyes. Cells (1) and (2) both would be in group 4 since they responded about equally to the two eyes. Cell (3) responded only when both eyes were stimulated; we can say only that it was not a group 1 or a group 7 cell.

One of the three did not respond at all to either eye alone, and thus its presence would have gone undetected had we not stimulated the two eyes together. Many cells show little or no synergistic effect; they respond about the same way to both eyes together as to either eye alone. A special class of binocular cells, wired up so as to respond specifically to near or far objects, will be taken up separately when we come to discuss stereopsis, in Chapter 7. These hookups from single cells to the two eyes illustrate once more the high degree of specificity of connections in the brain. As if it were not remarkable enough that a cell can be so connected as to respond to only one line orientation and one movement direction, we now learn that the connections are laid down in duplicate copies, one from each eye. And as if that were not remarkable enough, most of the connections, as we will see in Chapter 9, seem to be wired up and ready to go at birth.