Vision and the Brain

by David H. Hubel

Let me begin by stating that we all know the brain is complex. This statement is so obvious as to be almost insulting. But why do we think it is complex?

I can think of several reasons. One is that the brain does many complex things, or so we like to think. We walk, we talk, we perceive, we play the piano or the violin, and we do things like cough, vomit, and sneeze. Another reason is that the brain contains 10^{11} cells—give or take a factor of 10-literally an astronomical number. But anyone who has seen the size of a human liver and has looked at any part of it under a microscope will probably guess that it too has 10^{11} cells. There might be five basic types of cells in the liver, with each cell of a given type doing more or less the same thing. Nevertheless even the most committed hepatic physiologist would never suggest that the liver is more complicated than the brain. The number of cells obviously doesn't tell the story. If on the other hand you are aware that each brain cell, on the average, transmits information in the form of nerve impulses to maybe a thousand other cells and at the same time can receive information from about as many others, then the number indicating complexity goes up very quickly. Now we are getting into very large astronomical numbers, maybe even cosmic. It is very hard to think about such a complex structure, and our terms for summing up what it does, like perception and consciousness, are woefully inadequate. But it won't help us to complain about the complexity: we simply have to dig in, select one part to study, and see where that gets us.

I am aware that the basic topic at this colloquium is the future of science, and I was supposed to talk about future research on the brain. That is a little hard to do if one doesn't have a basis to go on. I think that even this group of scientists is not fully aware

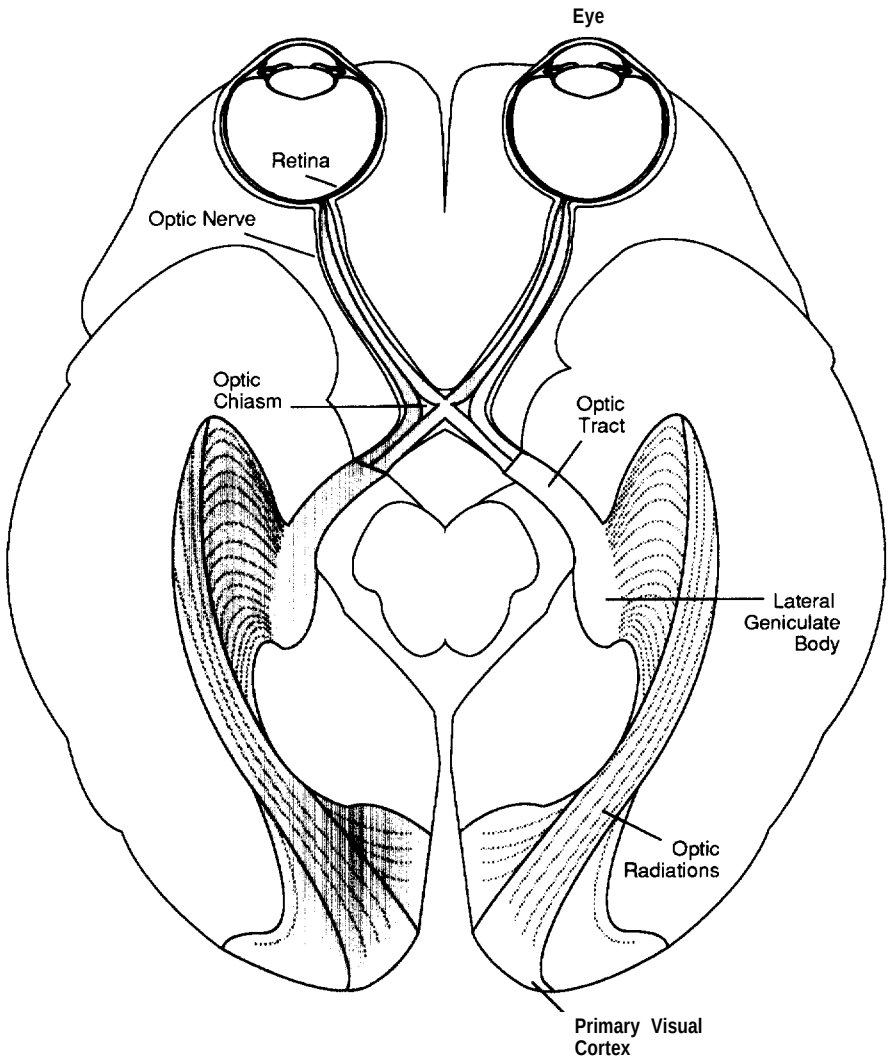

THE VISUAL PATHWAY

Fig. 1. The human brain and eyes seen from below. About 1 million optic nerve fibers come from each eye. At the optic chlasm half the fibers from each eye cross to the opposite hemisphere and travel back on the optic tracts to the lateral genlculate bodies. There the information is relayed to lateral geniculate cells, whose fibers pass back through the brain, in the optic radiations, to the primary visual (or striate) cortex. Note that because of this pattern of wiring, each hemisphere of the brain gets input from both eyes, and s given hemisphere, say the left, gets input from the two left half retinas, and consequently the right half of the visual world, from both eyes.

of what is going on in the rather arcane fields of brain neurophysiology and neuroanatomy. So instead of speculating about the future, I want to tell you some of what we've learned about the part of the brain concerned with vision. We know more about vision in the mammalian brain than about any other aspect of the central nervous function. The topic is a rich one, and even to give you a rough idea about it I will have to go into some technical details. But the main thing that I want to convey is a flavor for the sort of research that is going on now and for the sort of concrete facts we are learning about visual perception.

The mammalian visual system is remarkably similar among the different primates, so although the work I shall describe has been done on the macaque monkey and the squirrel monkey, the results apply almost unchanged to the visual system of the human brain. Let's begin by looking at the layout of the visual system from the eyes to the primary visual cortex at the back of the brain.

In Fig. 1 we are looking at the human brain from below. At the top of the figure are the two eyes out of the back of which come the optic nerves. Behind the lens of each eye is the retina, which contains a mosaic of 125 million light detectors called rods and cones. These light receptors make synaptic contact with other nerve cells in the retina; that is, their nerve fibers, or axens, split into a few or many branches that end on a second set of cells. These in turn have branches that end on a third set of cells called the retinal ganglion cells. The axons of these ganglion cells bundle together to form the optic nerves. About a million optic nerve fibers extend out from each eye. Some of the fibers stay on the same side of the brain, and some of them cross onto the other side. All their terminals end up in one of the two nests of cells known as the lateral geniculate bodies—geniculates for short-each containing roughly 1.5 million cells. The geniculates are deep in the head roughly between your two ears. They have a rather simple structure, in that any particular geniculate cell gets input from the optic nerve fibers and sends its output through the substance of the brain to what is called the primary visual cortex. or striate cortex, located at the back of the brain. Axons of the cells in the primary visual cortex project to a neighboring area, which then feeds into two or three other areas, and so on. Figure 2 shows a diagram of the visual pathway, which is actually made up of many millions of nerve cells.

Over the last twenty years people have come to understand fairly well how individual nerve cells work (Fig. 3). Without going into any of the rich detail, let me just say that one cell sends messages to another by events called nerve impulses. Whether a given cell fires or not depends on the sum of what it's told to do by other cells, some of whose impulses excite the cell in question, others of which inhibit it. The inhibitory influences are very important, as you will see.

Now let's consider what kinds of messages are sent through the visual pathway once light reaches the retina. It has been known since 1950 that a great deal of processing goes on between the light detectors and the optic nerves, and that the optic nerves carry rather sophisticated messages to the brain. Stephen Kuffler, the person most responsible for working this out, was my boss for some years. One of his favorite sayings, which fits the topic of all of these talks very well, is that the hardest thing to predict is the future. Now you can understand why 1 am keeping rather quiet about the future of brain research.

One of Kuffler's main contributions was to show that optic nerve fibers are carrying complex information of the following sort. Suppose I have an anesthetized animal facing a screen a couple of meters away and I put a microelectrode near or into one of the animal's optic nerve fibers. Then I shine lights

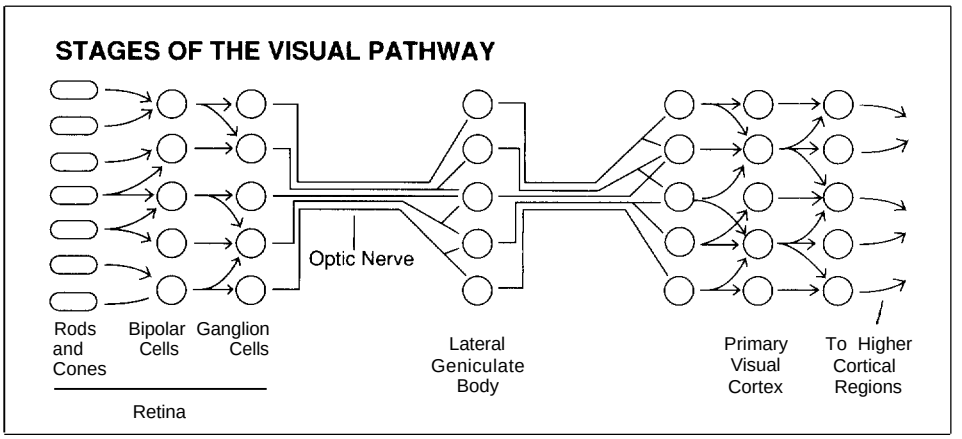

Fig. 2. Schematic illustration of the visual pathway from the retina to the higher cortical areas. At each stage one cell receives input from many cells at the preceding stage and passes information on to many cells at the next stage. Although only a small number of neurons are shown, each stage contains millions of neurons.

on the screen and ask what type of signal must reach the retina to influence this particular fiber (Fig. 4). It turns out that the area of retina influencing a typical optic nerve fiber is limited in extent, typically a circle about 1 millimeter in diameter. This region is called the receptive field of the cell. If we confine the light reaching this receptive field to a very small central region, we can drive each ganglion cell to produce up to 50 or so impulses during the onset (first tenth of a second) of stimulation. The cell then continues to fire at an average rate of up to 100 times per second. On the other hand, if we illuminate the whole receptive field, the cell responds at a much slower rate. Many of the optic nerve fibers hardly fire at all if you bathe the whole screen in light. That means that the ganglion cells are not interested in the amount of light hitting the retina but rather are interested in contrast. Each cell is making a spatial comparison between the amount of light hitting one tiny central region of its receptive field and the amount falling on the immediate surround. Illumination of the center excites the cell, and illumination of the surround inhibits it. Consequently these ganglion cells are described as having concentric "on'' -center and "off'' -surround receptive fields. Actually, there are two kinds of center-surround cells, those with "on" centers and "oft" surrounds and those with "off" centers and "on" surrounds. The first respond best to light spots on a dark background, and the latter to dark spots on a light background (Fig. 5).

If we now do a similar experiment with the cells in the lateral geniculate body we find that they respond in roughly the same way. Each individual cell takes care of a small region of retina, and the way that region influences the geniculate cell may again be described as a concentric center-surround receptive field. (Notice that the term receptive field refers to both the region of the retina influencing a given cell as well as the nature of that influence.) Thus the kind of information that the geniculate cells send to the primary visual cortex is of a very special kind. I should point out that hundreds or thousands of retinal ganglion cells and hundreds or thousands of geniculate cells are taking care of one small region of the retina. That is, these cells have receptive fields whose centers overlap within this one small retinal region. So one small spot on the screen will activate thousands and thousands of cells.

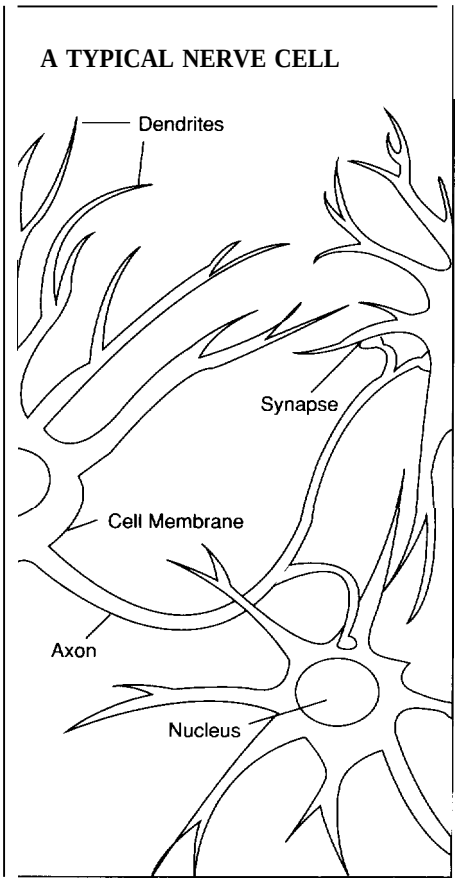

Fig. 3. A typical nerve cell may have from a few to over a thousand branches called dendrites, which receive signals from other nerve cells. Depending on the sum of the inputs (excitatory and inhibitory) the cell body fires or does not; that is, it sends out an electrical impulse that travels down the axon at a speed of about 10 meters per second. A nerve cell that fires rapidly does so at roughly the rate of a machine gun (an average of 15 times per second). When the impulse reaches the terminals of the axon, chemical messengers called neurotransmitters are released that provide input to the next set of cells. The relay of signals from one cell to the next takes place at specialized sites of contact known as synapses.

When we reach the primary visual cortex we find several stages of information processing. The cells at the earliest stage work roughly like the geniculate cells: they have center-surround receptive fields. At the second stage there are cells whose receptive fields are similar to the center-surround cells, in having an excitatory region and an inhibitory region. The geometries of these receptive fields, however, are different. They are designed to see either a light line on a dark background, a dark line on a light background, or an edge between dark and light. Moreover the line must have a certain orientation. The various receptive fields of these so-called simple cells and the response of one of them are shown in Fig. 6.

An engineer can give you a perfectly plausible diagram for how to get from a center-surround cell to an orientation-selective one. Depending on the engineer the diagram may differ, but the simplest circuit is to imagine a master cell getting input from a lot of center-surround cells, each of which differs in the positions of their receptive-field centers. If these receptive field centers are arrayed along a line, then the master cell gives a maximum response when the stimulus covers all of the center and only a small part of the surround of each of these ancestral cells, as shown in Fig. 7. A line of light made by a slit does the trick best of all. The precise circuit for this sort of thing is not known, and I am not going to discuss the possible circuits any more here.

Beyond the second stage cells work

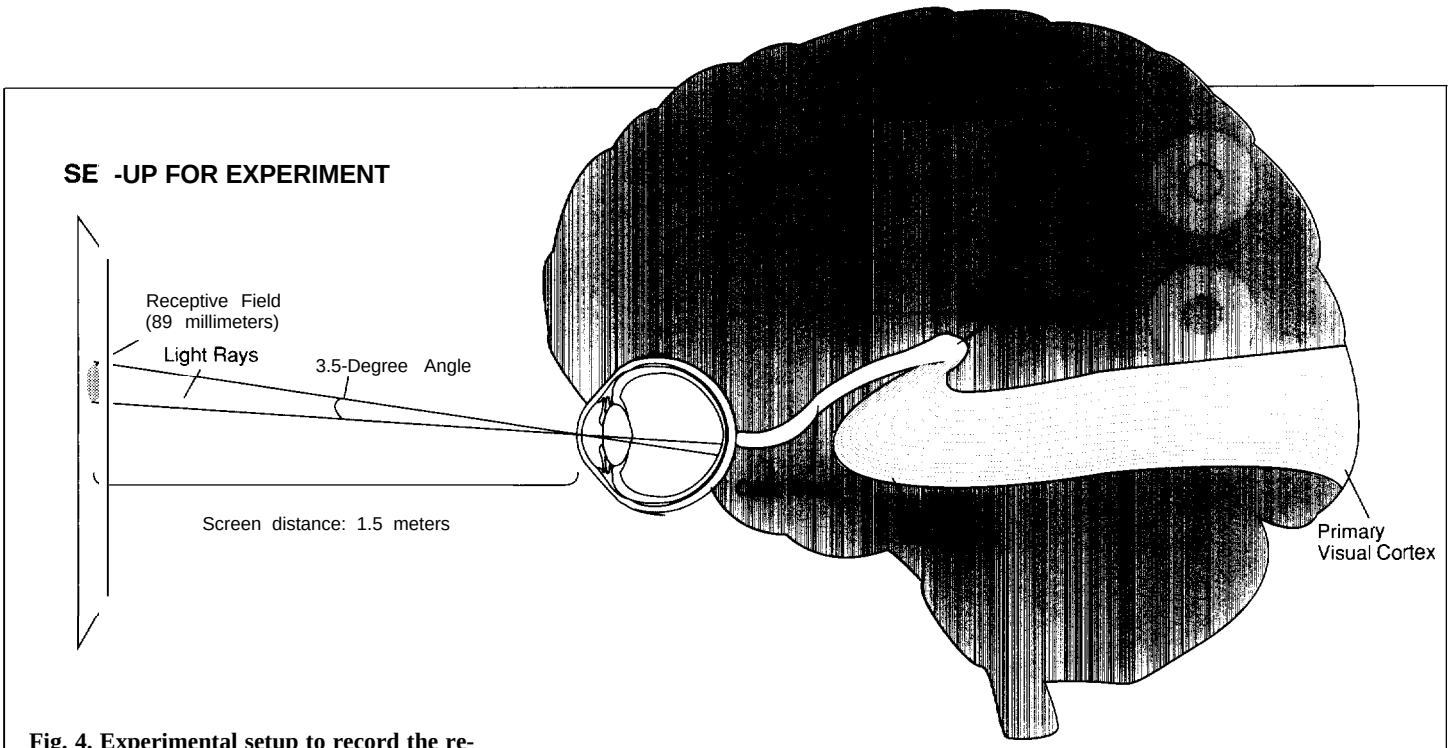

Fig. 4. Experimental setup to record the response of a single nerve cell to visual stimuli. An anesthetized animal is facing a screen about 4.5 feet away. Stimuli projected onto the screen are focused by the lens of the eye onto the retina. From there signals travel along the visual pathway (seen here from the side).

The receptive field of a single ganglion cell in the retina is a circular area about 1 millimeter in diameter, which subtends an angle of about 3.5 degrees. At a distance of 1.5 meters, a distance of 89 millimeters on the screen corresponds to 1 millimeter on the retina.

in even more complicated ways. Just as before, any given cell takes care of a small region in the retina corresponding to a small region of its visual field, for example, a portion of the screen in the experiments I described above. But these so-called complex cells don't respond to bright small spots anywhere in their receptive fields. What they like is a line moving across the region. They are also very fussy about the orientation of the line. Most of the cells are so fussy that if you change the direction of the line by more than 20 to 30 degrees they don't respond at all. Again, by "line" I mean either a light line on a dark background, a dark line on a light background, or an edge or boundary, say, between dark and light. Some cells are finicky about which of these three they respond to; others are less so. Some of them are also fussy about the direction in which the line is moving.

Now let me describe in a little more detail the experiments that demonstrate the response of these complex cells. As usual, we have an anesthetized animal facing the screen and an electrode poked into a single complex cell in the primary visual cortex. The electrode is connected to an audio monitor and each impulse from the cell is recorded as a click. The stimulus is a straight line of bright light on a screen. We project light through a slit to create this pattern of illumination, and we sweep the slit across the screen in a direction perpendicular to its orientation and repeat this for all orientations—as if we were painting light on the screen. In this way we are able to map out three things about the cell's receptive field. First we map out the receptive field of the cell, the area of the screen over which the cell can be influenced. When the line of light is moving over the cell's receptive field, the sound from the audio monitor becomes a din of clicks—the nerve cell is firing at machine-gun speed—but as the line goes past the region the sound dies down to a few isolated clicks and then to silence. As we change the orientation of the line between vertical and horizontal, the intensity of the sound corresponding to the number of impulses per second clearly varies, so

continued on page 24

CENTER-SURROUND CELLS

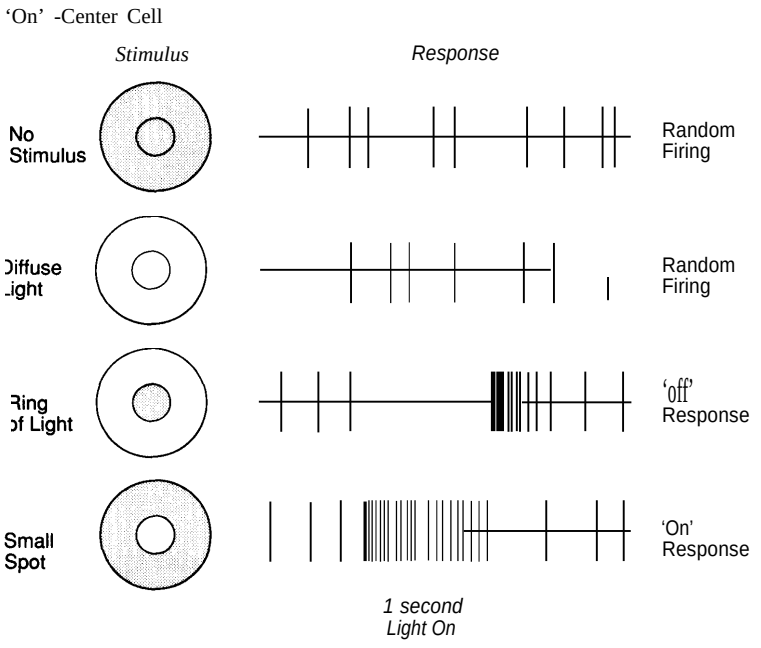

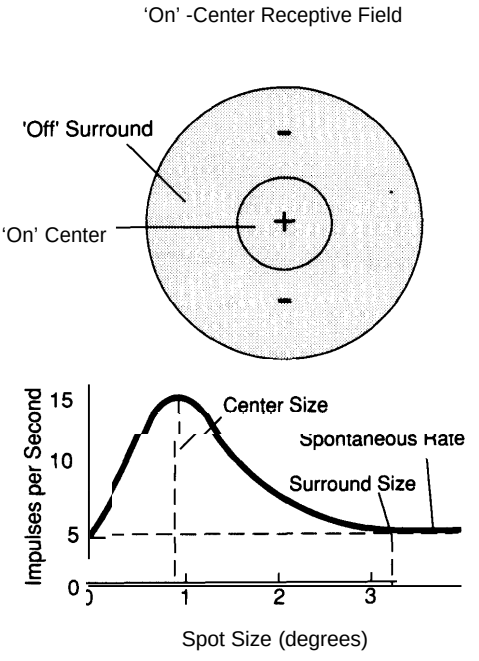

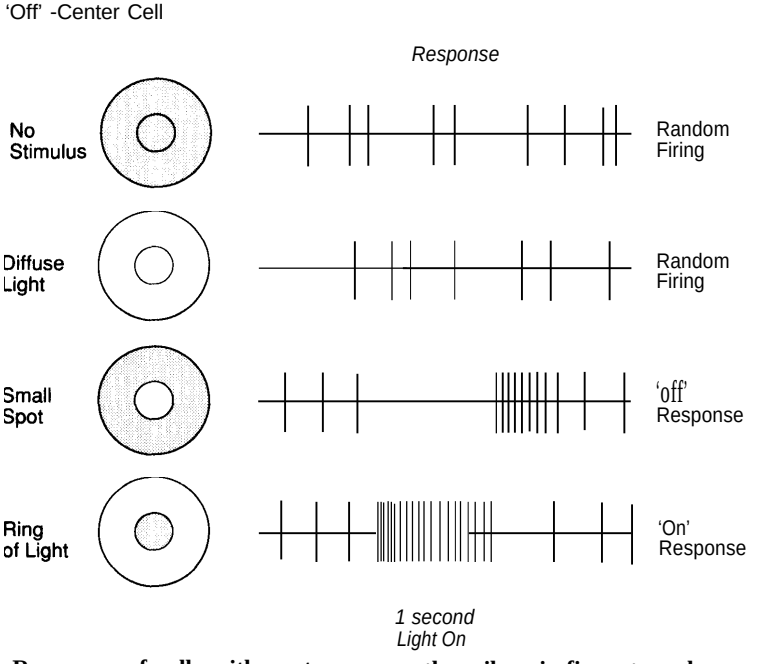

The graph shows the response of a single "On''-center geniculate cell versus spot size. The spot changes from one that is smaller than the center to one that covers the entire receptive field. The maximum response occurs for a spot that covers the "On" center only (1 degree). As the spot grows larger and invades the inhibitory region the number of impulses decreases until the spot size is 3 degrees (the size of the entire receptive field). Further increase in spot size does not change the cell's response.

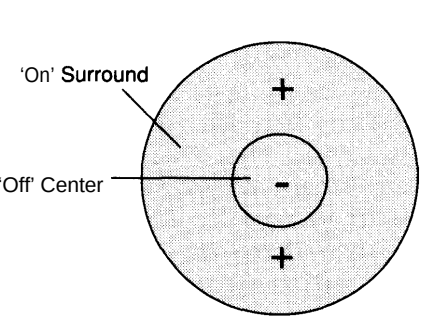

Fig. 5. Responses of cells with canter-surround receptive fields to various stimuli. The recordings are from genicuiate ceiis, which have the same pattern of response but fire much more siowiy than retinai gangiion ceiis. The response ilne is 2.5 seconds iong. Each verticai line represents a single Impulse In the genicuiate ceii being recorded. With no stimulus the ceil fires slowiy (a few times par second) and more or iess at random. With diffuse light covering the entire receptive field, the ceil again fires at random and just a iittie faster. With a smaii bright spot covering just the center, the ceil fires rapidiy. A ring of iight covering the surround stops the ceii from firing. When the stimulus is turned off the celi produces a brisk burst of impulses lasting about a second. The receptive fieid Is thus described as an excitatory ("On") center and an inhibitory ("Off") surround. The responses of the "Off" center, "On" surround ceiis are aiso shown.

Fig. 6. The receptive fields of simple cells in the cortex are designed so that the ceil responds to (a) dark lines on a light background, (b) light lines on a dark background, or (c) a straight edge between dark and light regions. Like the center- surround cells, the simple cells have excitatory and inhibitory domains, A small spot anywhere in the receptive field will give a small response, either inhibitory or excitatory depending on its location, but the maximum response is obtained by stimulating the entire excitatory region and neither of the inhibitory regions. The job is done best by a slit of light whose width is about 2 minutes of arc.

Five recordings from a cell designed to see a light line on a dark background. This particular cell exhibits no spontaneous activity (some do) and also does not fire if the whole field is illuminated. When a light line fills the excitatory region the cell produces a maximum "On" response. If the light line is moved to the inhibitory region, the cell fires after the stimulus is removed. If the line covers only a small part of the excitatory region and a proportionally small part of the inhibitory region, the cell again fails to fire.

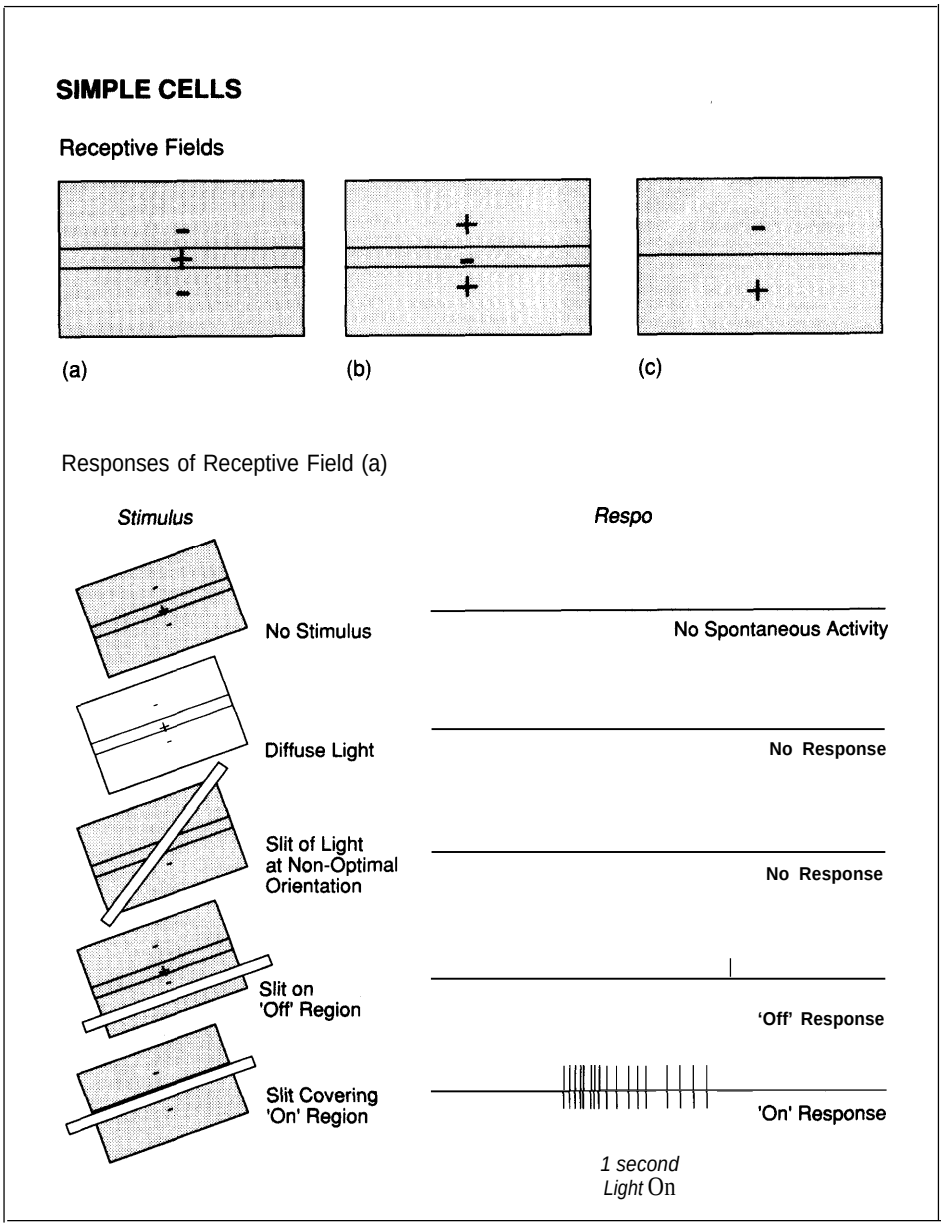

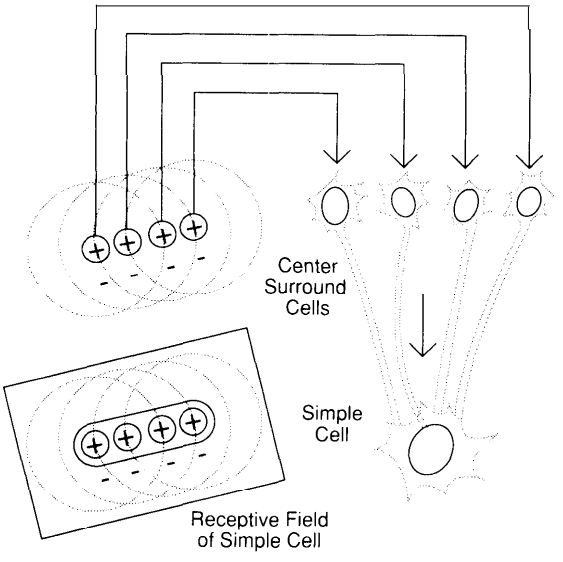

"SIMPLE" ORIENTATION-SELECTIVE CELLS

Fig. 7. One possible circuit for converting a center-surround response to an orientation-selective response. The orientation-selective cell receives input from a series of center-surround cells whose centers overlap and are arrayed along a line. When summed together, the center-surround receptive fields have an excitatory region that looks like a long narrow rectangle, flanked by inhibitory regions on either side. The receptive field looks very much like (a) in the preceding figure.

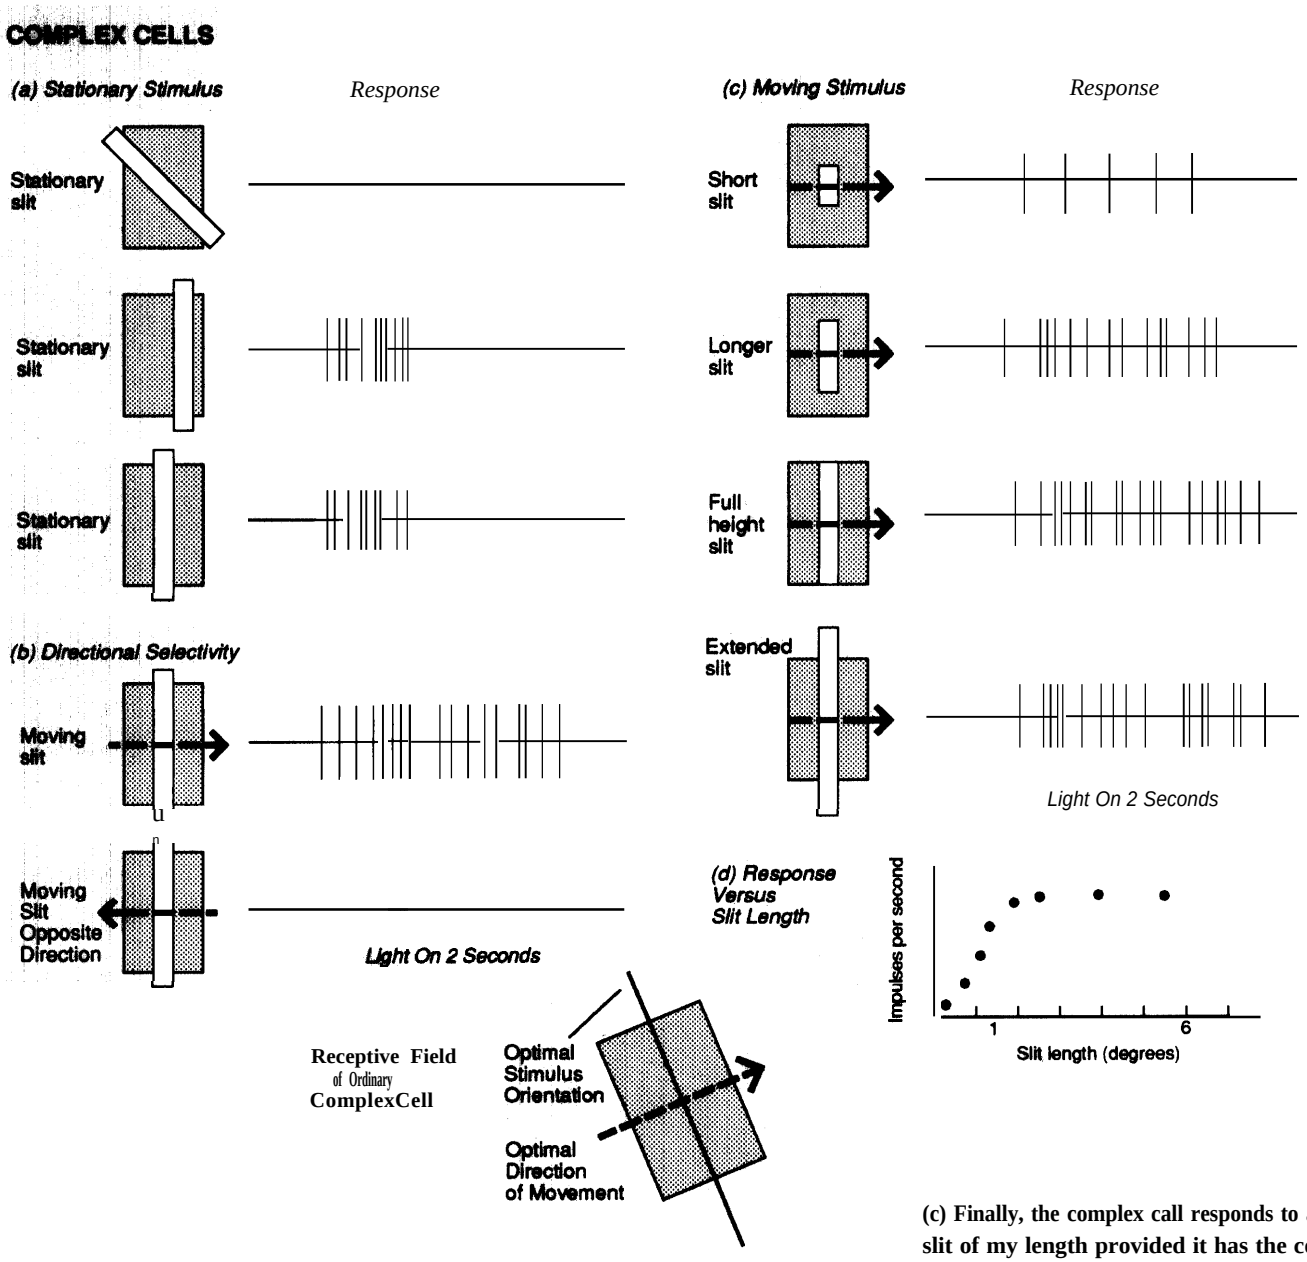

Fig. 8. A complex caii in the cortex generaiizes the response of a slmpie cell. (a) The complex cell responds to a long narrow silt of light placed anywhere in its receptive field (rectangle) provided the orientation of the line is correct. Changing the orientation by ± 10 degrees will produce weaker response; a larger rotation gives no response at ail. The optimai width of the siit is 2 minutes of arc, the same as for the simple ceii.

(b) The complex ceii responds best to a siit of light moving perpendicuiar to its direction of orientation. For about 20 percent of these complex ceils, the siit must move in one direction oniy. Motion in opposite direction faiis to produce a response.

(c) Finally, the complex call responds to a light slit of my length provided it has the correct orientation but the maximum response is obtained when the siit is at ieast iong enough to cross the receptive fieid (2 degrees).

(d) The graph shows the response versus siit iength for an ordinary compiex ceii. The response increases untii the siit iength is 2 degrees (the width of the receptive fieid) and then remains constant as the siit iength is increased still more.

END-STOPPED CELLS

(a) Moving Stimulus of Varying Length

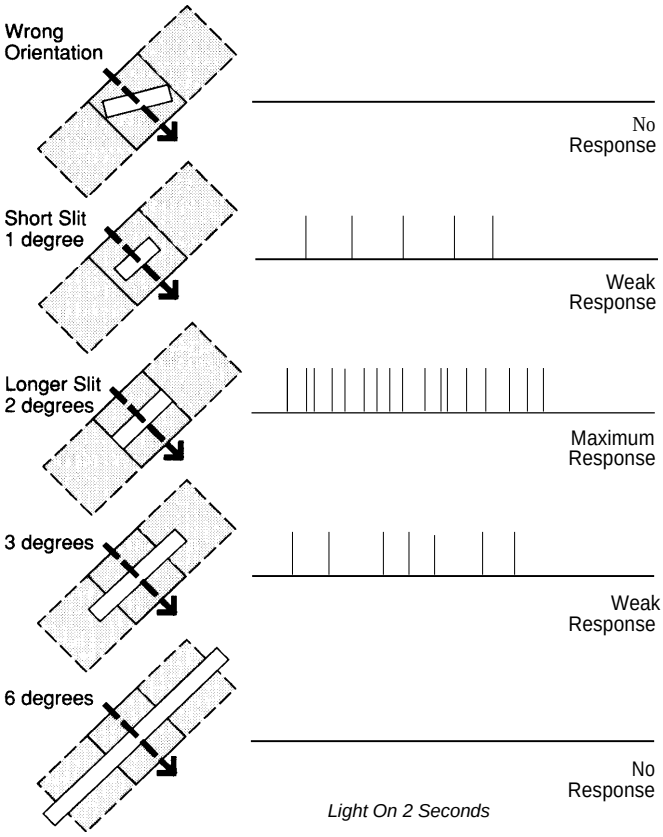

(b) Orientation of Inhibitory Region

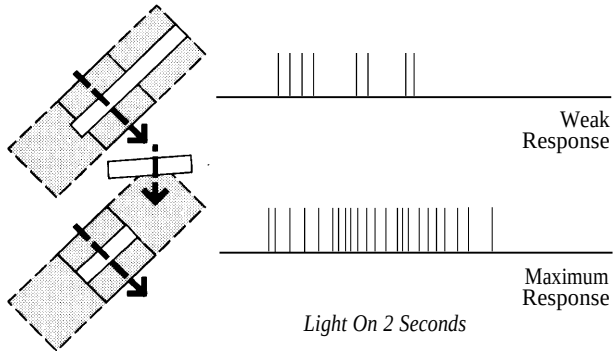

Receptive Field of End-Stopped Cells

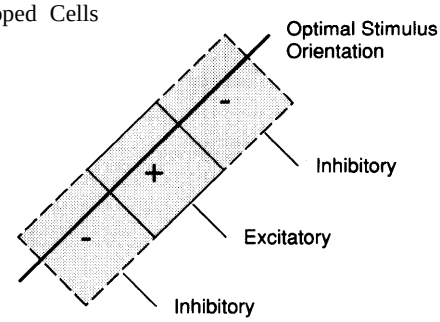

Response versus Slit Length

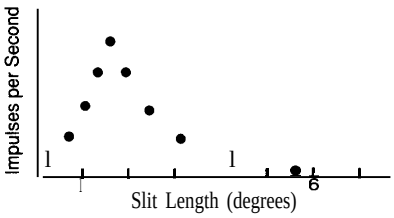

Fig. 9. (a) End-stopped cells, like complex cells, respond to lines with a specific orientation moving perpendicular to that orientation and as the line gets longer the response gets stronger. But unlike what one finds with complex cells, this cell shows a weaker response when the line is longer than about 2 degrees. At a line length of 6 degrees, the cell gives no response at all from the excitatory region. Evidently the (activating) excitatory region is flanked by inhibitory ones on either side as shown in the diagram of the receptive field. Each record has a 2 second duration. The graph summarizes the results of varying the slit length. (b) These records show that the optimal orientation for the inhibitory region is the same as for the excitatory region. In the second record a nonoptimally oriented slit in the inhibitory region does not reduce the response to an optimally oriented slit of optimal length in the excitatory region.

continued frorn page 19

we can map out a preferred orientation for the line as well as a range of orientations over which the cell gives some response. Finally we swing the light back and forth, and the dramatic alternation between rapid firing and silence tells us that the cell responds to motion in one particular direction and not to motion in the opposite direction. The experiment is so clear that we made a movie of the stimulus on the screen and simultaneously recorded the audio signal generated by the cell's response. In the movie we draw in the receptive field, the preferred orientation of the line, and the preferred direction of motion of the line as these become evident from the audio response. Some results of the experiments on directionally selective complex cells are shown in Fig. 8.

We made a similar movie for another type of complex cell that is fussy about the length of the line. We call that an "end-stopped" cell. The cell responds very well to a short line but very badly to a long line. Apparently inhibition plays an important role in the functioning of end-stopped cells, just as it does in the functioning of center-surround cells. The receptive fields of these cells extend beyond the region where you get a big response. But the only way you can know that is to make the line longer and find that then you don't get any response at all from the cell. Evidently stimulating the region beyond the short line has the effect of inhibiting the cell, and if you inhibit the cell as much as you excite it, then it just sits there and does nothing. Figure 9 shows results of an experiment as well as a diagram of the end-stopped cell's receptive field.

Now let me add one more thing to try to get at why we think the Almighty would have given us end-stopped cells. Suppose you sit back in your chairs and look at the form in Fig. 10. If you fix your gaze on a point toward the center of the form, millions of complex cells

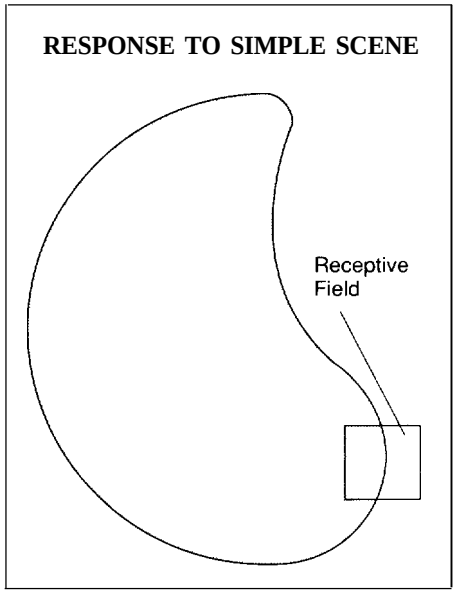

Fig. 10. Diagram to illustrate the effect on typical cortical cells of a simple scene, a kidney-shaped dark patch on a white background. If we fix our gaze on a point toward the center, we can imagine that the only cortical cells to be affected will be ones whose receptive fields are cut by the boundary, and whose optimal stimulus orientation happens to be appropriate. For example, a cell whose receptive field is vertically oriented and is stimulated by the boxed region in the figure will be activated. A cell whose field is entirely inside or outside the blob will be unaffected by the stimulus.

in your brain are being activated by the borders. In fact the only cells that are going to be tied up by this stimulus are the cells that are mediating the borders. It turns out that as you go farther and farther toward the center of the form, things are arranged in such ingenious ways that the number of cells required to give information about the interior becomes less and less. If you only consider end-stopped cells, the number of cells that are tied up by this stimulus is rather small. Only the end-stopped cells whose receptive fields happen to coincide with the regions of high curvature will respond. (Remember high curvature is essentially equivalent to small line segments in a very small region. Figure 11 shows how the receptive field is designed to respond to curves.) These phenomena are very counter-intuitive. You wouldn't think that your vivid impression of this homogeneous form would be conveyed at some stage of the brain by cells that aren't even telling you anything about the interior. It is the fact that information is coming from the borders and no additional contradictory information is coming from the inside to tell you that the contrast has changed, that lets you know that the whole form is filled. At first sight you may find this a hard pill to swallow, but it happens to be the way the brain works. If an engineer wants to build an image-processing device, he would probably invent a very similar design. He has to pay for all the transistors that take care of the innards,

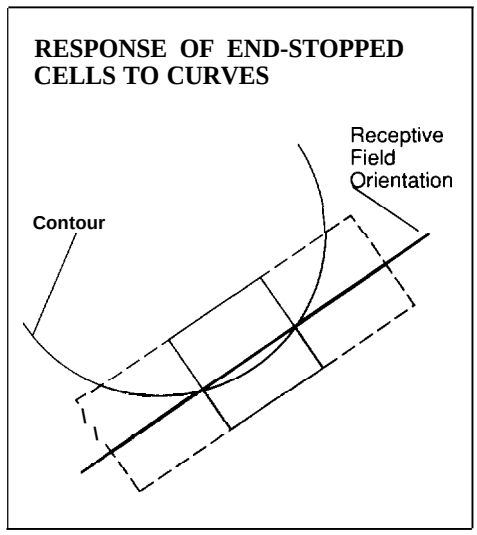

Fig. 11. This diagram shows how end-stopped cells are well designed to respond to curves. The segment of the curve passing through the excitatory region is oriented to cause a large response, whereas the segments through the inhibitory regions have the wrong orientation to inhibit the response. The sum total will thus be a large response.

and so he wants a machine that uses as few of them as possible.

So far we have discussed cells that distinguish contrast, or differences in brightness, and are therefore involved, loosely speaking, in the perception of form. The subject that I want to discuss during the remainder of this talk concerns our intuition that visual perception is not a unitary thing but must be subdivided. When we look at a scene, we are not necessarily conscious of the various subdivisions of perception. But suppose you ask an average person to break up vision into its parts. Most people would probably say you have form, color, movement, depth, maybe texture, and a few others. If you gave the right cues, even a Boston taxi driver would give you some list like this, so you do not need a scientific or neurobiological background to come up with it. This subdivision is intuitively reasonable to us. Now it turns out—and it didn't have to be so-that the brain divides up vision pretty much according to this list.

Let's take a look at the anatomy of the lateral geniculate body because that is the first place where the division of the visual pathway is rather obvious. Figure 12 is a cross section through the lateral geniculate body. From one side to the other is about 3 or 4 millimeters, and because of the stain that was used each dot is a cell body. This picture alone tells you a great deal about the geniculate's structure. It is rather like a tiny jelly roll consisting of a number of layers, one topping the other and all rolled up in a sort of curved way so that each cross section parallel to the one in Fig. 12 would look very much the same. Each layer in the geniculate receives input from either the right eye or the left. Thus, in the right geniculate body shown here, all the cells in the top layer get their input from the right eye, all cells in the next layer get their input from the left eye. The whole sequence of inputs from top to bottom goes right, left, right, left, left, right. Why the order changes near the bottom nobody knows; it may just have been to make it hard for us to remember.

The main feature that I want you to notice is the obvious difference between the two ventral (underneath) layers and the remaining four dorsal (upper) ones. You can see that the cells in the two ventral layers are bigger and more thinly scattered. If you looked at this cross section with more powerful methods you would see other differences. It has been clear for a century that the geniculate is subdivided into these two distinct regions; the ventral layers are called magnocellular because the cells are big, and the four dorsal layers are called the parvocellular layers, "parve" for small. Moreover these two kinds of cells get their inputs from two different kinds of cells in the retina. The magno get their inputs from big retinal cells, and the parvo from small ones. At later stages in the cortex, these two branches of the visual pathway, magno and parve, don't merge but keep their separateness and seem to have different functions. Although both magno and parvo geniculate cells have receptive fields with a center-

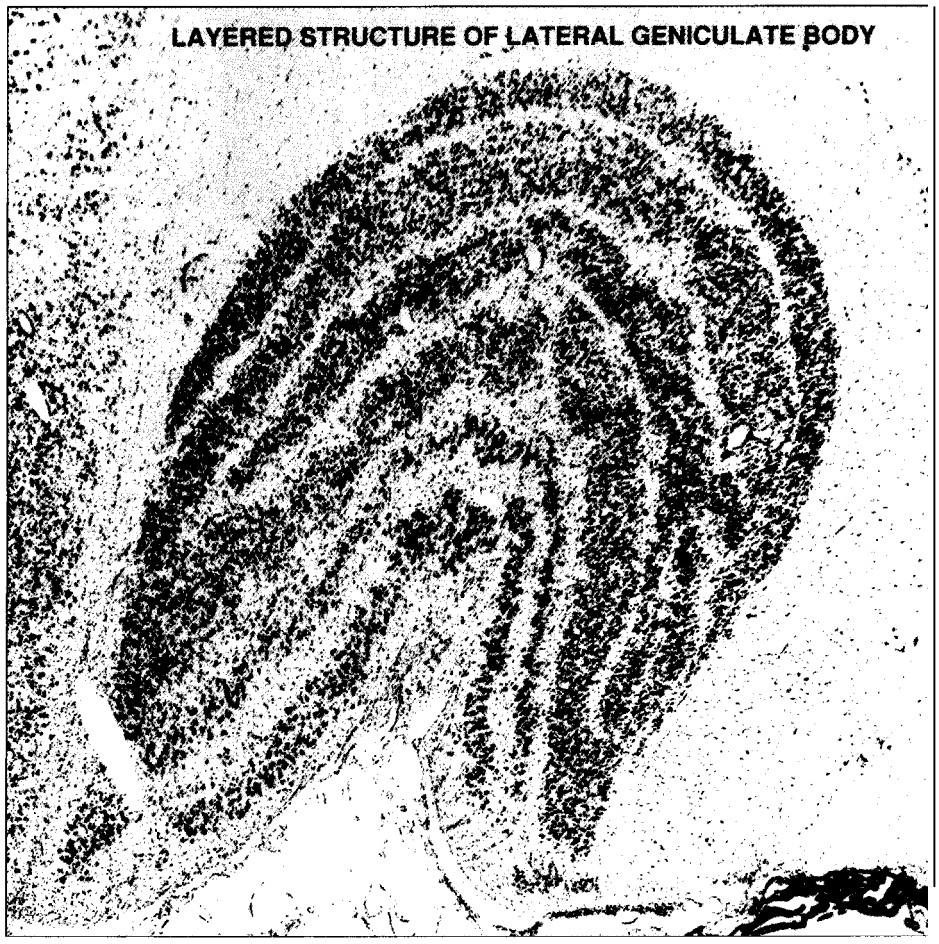

Fig. 12. Microphotograph of a cross section through the right lateral geniculate body of a macaque monkey. Each dot is a cell body, stained with cresyl violet dye. Each of the six Iayers gets input from one eye only. A human lateral geniculate body would look almost identical to this.

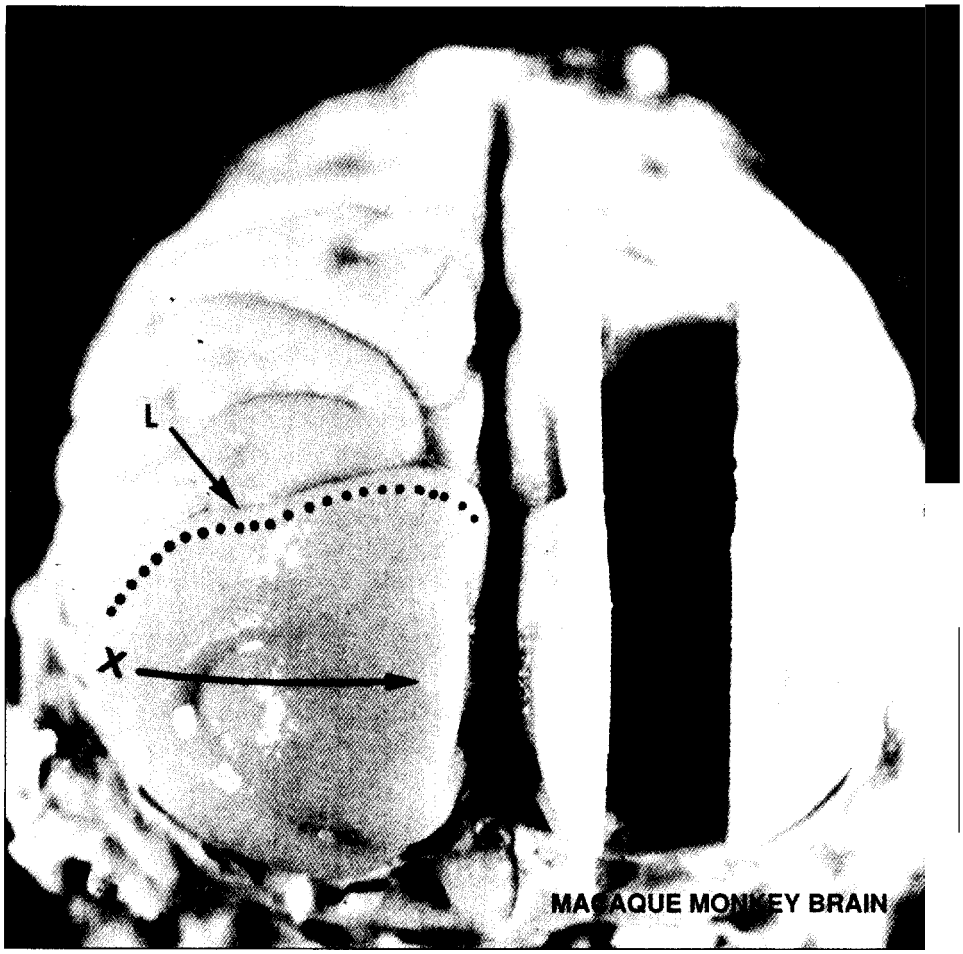

Fig. 13. A macaque monkey brain viewed from behind. The region in the foreground (below the dotted line) is the primary visual cortex, the part that is exposed on the surface. L is a deep cortical fissure (the lunate sulcus). X is the representation of the center of gaze. As we proceed from X in the direction of the arrow, the part of the visual field mapped goes to the right along the horizon.

surround target-like arrangement, the receptive field centers of the magno cells tend to be bigger than the field centers of the parvo cells. It is as though each magno cell got its input in parallel from all three kinds of color cones and are therefore color blind. The parvo cells, on the other hand, are very strongly color sensitive. They respond to color as though their centers got their input from one color cone only and the surround from one of the other two color cones only. A second difference concerns sensitivity to small changes in luminous intensity: the magno cells are much more sensitive than the parve. When the receptive-field center is just 5 percent brighter than the surround, the magno cells respond very well, whereas the parvo cells won't respond until the intensity difference between center and surround is at least 20 percent.

Now let's see what happens to these two pathways in the visual cortex. To do so we need to look at the anatomy of the cortex more closely. Figure 13 shows what a macaque monkey brain looks like if you remove the top of the skull. The primary visual cortex (also called visual area 1) occupies most of the area below the dotted line, but part of it is tucked underneath in a slightly complicated way. Its area is about the size of a credit card, and it has a thickness of two millimeters—thicker than the average credit card, unless you have the gold kind.

If you cut out a chunk of the cortex (as in Fig. 13) on the right side and walk into the hole and turn left, you would see the cross-section of the primary visual cortex shown in Fig. 14. The richly layered structure is made visible by an appropriate stain. The axons of the geniculate cells come up vertically through the lower layers of the cortex, branch again and again, and finally terminate in layer 4C, the very dark layer a bit more than halfway down. The 4C cells send their output to the upper layers of the cortex, and the upper layers send their outputs to other regions of the brain as well as to other layers in the primary visual cortex, We will be particularly interested in the projection to visual area 2, which borders visual area 1 (that is, it is right above the dotted line in Fig. 13), The projection is done in a very precise way, so that a tiny region in visual area 1 will send its output to a tiny region in visual area 2.

I said earlier that the two branches of the visual pathway, parvocellular and magnocellular, maintain their separateness in the cortex. In particular the magnocellular layers of the geniculate transmit impulses to the top half of ayer 4C (4Cα), which in turn transmits its output to layer 4B. Layer 4B then sends *its* output to visual area 2, The parvocellular layers transmit impulses o the bottom half of 4C (4Cβ), which transmits its output to the deep part of layers 2 and 3. The output from layers 2 and 3 again goes to visual area 2. That's as much as I want to say about these two pathways until later when

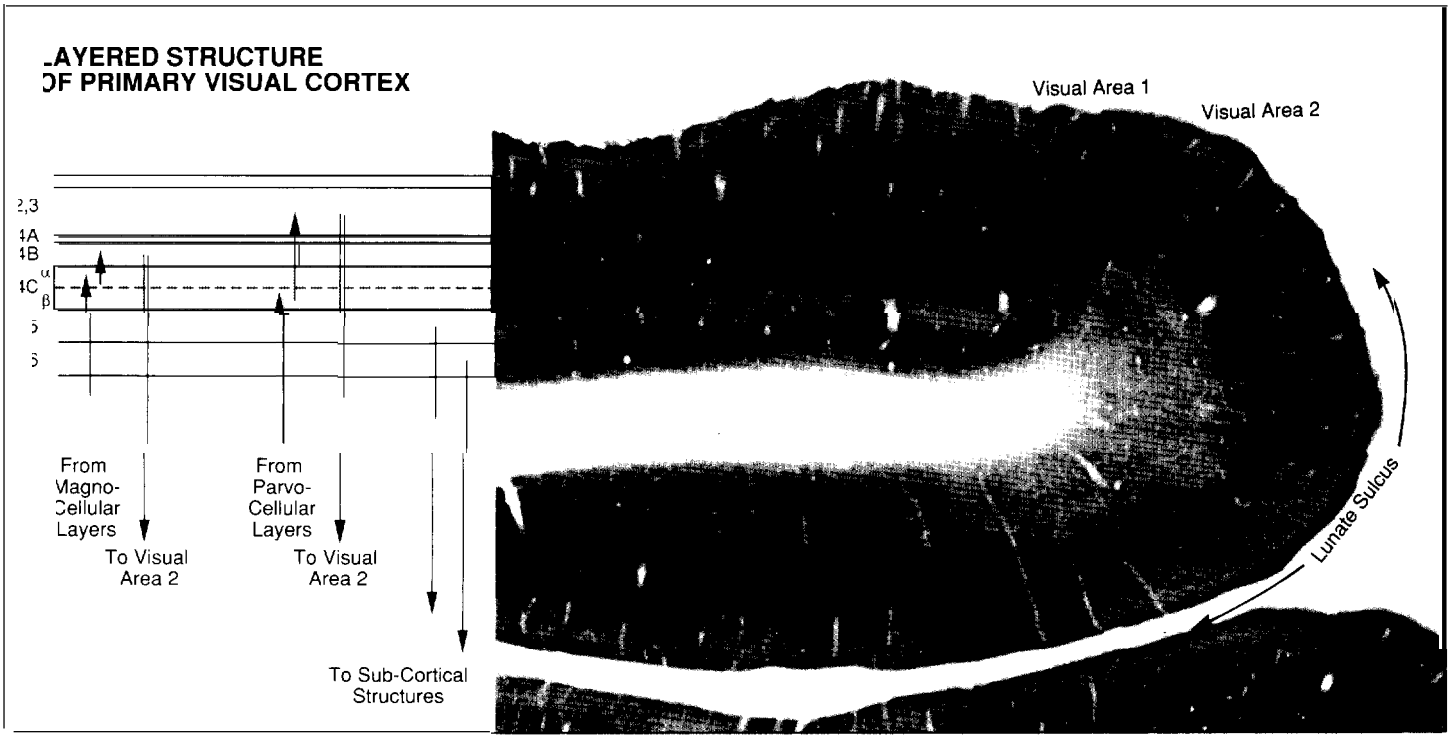

Fig. 14. Here we see a low-power cross section of the primary visual cortex, roughly what we would see if we were to walk into the cleft cut in the right hemisphere of Fig. 13 and look to our left. The cortical layers can be clearly seen. The pattern of layers changes as we go from the primary visual cortex to visual area 2. The transition between visual areas 1 and 2 occurs at the dotted line in Fig. 13. The deep fissure known as the lunate sulcus is visible in Fig. 14 just to the right of the transition between visual areas 1 and 2. About a dozen blobs can be seen above the very deeply stained layer 4.

Diagram at left shows the input to the cortical layers from the lateral geniculate body and the output to other regions of the brain.

we discuss visual area 2. About 1978 we and others began to suspect that, at least at the cortical level, there is a third branch of the visual pathway. By using a stain for cytochrome oxidase we were able to distinguish periodic regions in the upper layers (layers 2 and 3) of the cortex that were staining darker (see Fig. 14). In Fig. 15, which is a view of the cortex from above, these darkly stained periodic areas look like rather punctate oval regions about one-half millimeter apart. We call them "blobs" because of their appearance.

Around 1979 Margaret Livingstone and I were able to record the responses from individual cells in the blobs by driving a microelectrode parallel to the upper layers of the cortex. We had thought that all the cells in the upper layers are like the two kinds of complex cells I described before, either ordinary complex or end-stopped. It turned out, however, that every time we got our electrode into a blob, we

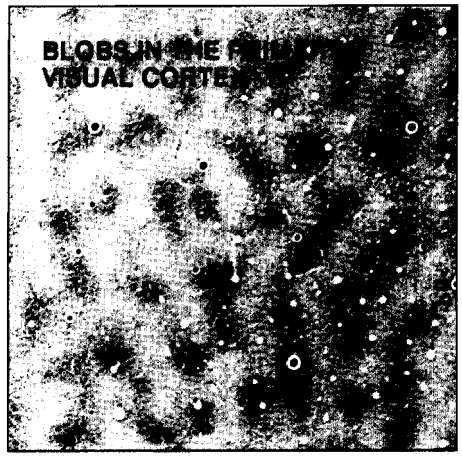

Fig. 15. Low-power picture of cytochrome oxidase blobs, in a piece of primary visual cortex cut in a plane parallel to the layers, above layer 4. We are viewing the blobs as if we were standing above the cortex in Fig. 14, looking down. Blobs are about 0.5 millimeter apart.

were likely to record five or six cells in a row that had absolutely no orientation specificity. They seemed to be of the center-surround style. About half of them ignored color, but the other half were richly involved in color and worked in a very specific way.

I want to discuss the variable color, so I will now describe how a color-sensitive blob cell works. It has a receptive field, with center and surround (Fig. 16). The center is likely to get its input from two types of cones in opposition, so that, for example, illuminating red cones has the effect of exciting the center and illuminating green cones inhibits it. If you happen to stimulate the center of the receptive field, the cell will turn on or off, depending on whether the light is red or green. If you use white light the inputs counteract each other because white light contains both long and middle wavelengths. The cell just doesn't respond at all to white light. The surround works in just the opposite direction, red inhibiting and green exciting. Now the ramifications of this

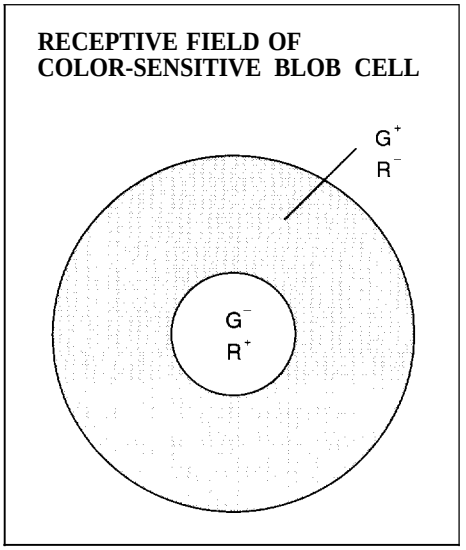

Fig. 16. Diagram of the receptive field of a red-on green-off double opponent cell. The cell's firing rate is increased by a small red spot and slowed by a small green spot. Large spots have no effect.

are many. For example a small red spot tells the cell to fire faster but if the spot is large, it has no effect. The only thing that will drive this cell to fire faster is a small red spot, a red edge, or a long red line that stimulates all of the center and very little of the surround.

I am not going to spell out all the implications of this, but one consequence is that our perception of color involves not only the wavelengths coming from the object we are looking at, but the difference between the wavelengths coming from that object and the wavelengths coming from other objects in the scene. Thus *space is* involved as well as *wavelength.* For many years Edwin Land has been presenting demonstrations aimed at convincing people that space is involved in color perception just as much as wavelengths. I will now give you a kind of watered down version of one of these demonstrations. It will at least help you to understand why a conservative Canadian-born would come to this meeting in such a garish, bright red tie.

Following in the footsteps of George Wald, I will take off my coat, but for this additional reason: to show you the full glory of this tie. Everyone I think would agree, except those who are frankly color blind, that this is a red tie. What I propose to do is first of all bathe myself in long wavelength light, that is, what we are used to thinking of as red—and see what the tie looks like. That is what I will do, so perhaps we can turn off all the light in the room, and I mean *all.* Now if we can have the red projector turned on, I think you will agree that if anything the tie appears to be a kind of anemic red. It is a pale ghost of what it was before, and yet the wavelengths that are coming to your eyes from the tie are just what they were before; they have not changed at all. If we now turn the red light out and turn on the short wavelength (blue) light, the tie looks very dark, naturally, because a red tie by definition is one that does not reflect back short wavelengths. So the tie isn't shining back much of anything to you. Next we can turn off the blue light and turn the red light back on, and you again see the anemic, pinkish, washed-out red. If I were now to add the short wavelengths, I think you know that nothing new is going to come from the tie. So let's add the blue. Now you see the tie and at least from where I stand it bursts forth again in all its glory. Yet what could be more counter-intuitive than that? This is so counter-intuitive that people made fun of Land for the first twenty years of his presentations, saying that he was using magic and things like that. But these are very real phenomena, not magic, unless you want to think of biology as magic and insist that only physics is real.

Why would the Almighty wire up our brain in such a way as this? I think the answer is reasonably simple. If we are out under the blue sky and look at a colored object and then come in here and look at it under tungsten light, our estimation of the color stays remarkably constant. To realize that this constancy in color perception is not a trivial thing, try going outside and taking a picture of a white shirt and then coming inside and rephotographing it under tungsten light. You get a pink shirt on the one hand and a white shirt on the other, or else you get a perfectly good white shirt under the tungsten and a blue shirt under the blue sky. The light source makes a big difference to the camera. The camera isn't equipped to factor out the light source. Our brains are so equipped, by some mechanism like the blob cells that compare wavelengths in different regions of their receptive fields. The result is that the perceived color remains the same despite changes in the light source, and our brain does the job so well that it is hard to convince ourselves that it really is solving a problem. The

VISUAL AREAS 1 AND 2

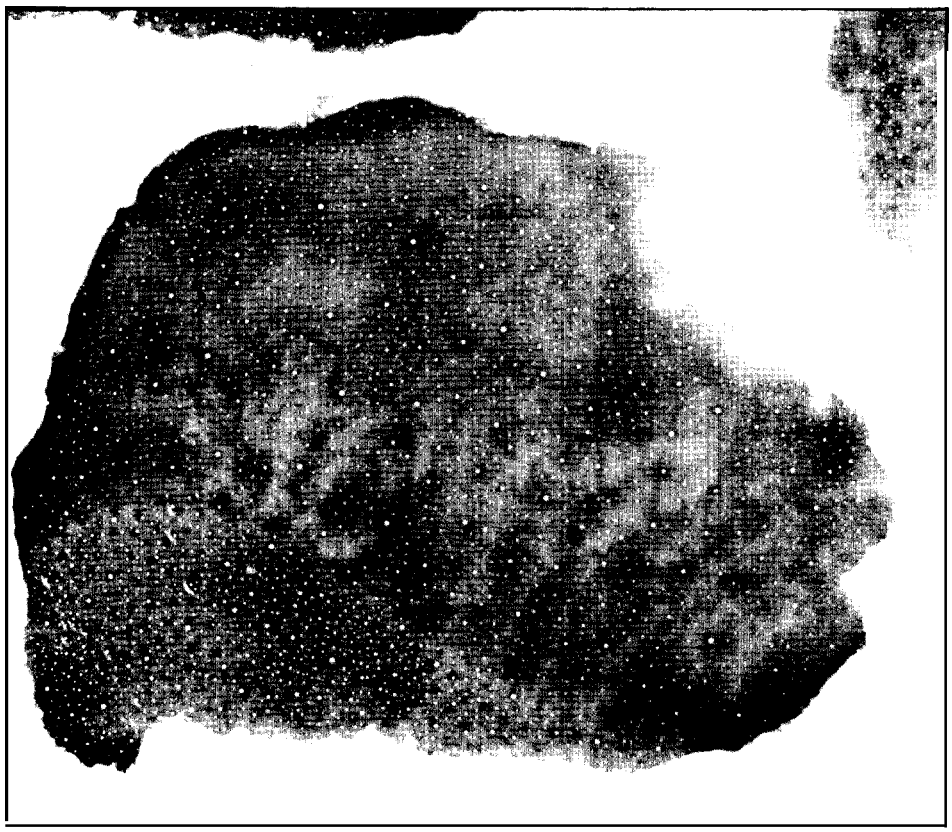

Fig. 17. A section parallel to the layers of the visual cortex of a squirrel monkey at much lower power than Fig. 15. To the lower left in the figure is the primary visual cortex, with blobs that are 0.5 millimeter apart. To the right is visual area 2, with thick, thin, and pale stripes running at right angles to the border between visual areas 1 and 2.

white is white. Why shouldn't white look white wherever you go?

It is indeed a complicated question, and many questions on perception have just that kind of complexity. *Exactly* how these blob cells solve the problem is not clear. One can go a certain way, but it is a complicated theoretical question and I think it hasn't been worked out satisfactorily so far.

The last topic that 1 want to talk about is visual area 2. If you stain visual area 2 for cytochrome oxidase and look at it, you don't see **blobs** but rather a pattern of dark stripes alternating with pale areas (Fig. 17). The stripes extend the full length of visual area 2, which is about 8 or 10 millimeters. and they appear every 4 or 5 millimeters. Furthermore the dark stripes are of two types, alternately thick and thin. This pattern gave us a hint that if we were to record from these regions we might find physiological differences. It also gave us a hint that the connections between visual areas 1 and 2 might have something to do with these blob and non-blob regions. We were able to show through anatomical work that the blobs project to the thin stripes and *only* to the thin stripes. I wish there were time to show you this convincingly, but it would take **a while and it is a bit technical.** We also found that many blobs project to a single thin stripe and inter-blob regions don't project to thin stripes at all. Further-more the traffic is two-way. Blobs connect to thin stripes, which connect back to the blobs. The inter-blob regions in visual area 1 project to the interstripes, the pale regions of visual area 2. And finally the thick stripes in visual area 2 get input from layer 4B in visual area 1, the terminus of the magnocellular system. To sum it up, the magnocellular branch of the visual pathway is represented in the thick dark stripes, the parvo-interblob branch by the pale stripes, and the blobs are represented by the thin dark stripes. We don't know the **source** of the input to the blobs, but it probably from magnocellular and parvocellular branches of the visual pathway. The three pathways have in some ways kept their separateness right up to visual area 2, and as I'll describe below, each one seems designed to perform a different function (Fig. 18).

When we record from visual area 2 we find that the cells in the thick stripes are very orientation-selective and very movement-sensitive, more so than the other regions. so we have reason to think that these cells are involved in perception of movement. They are also involved in stereoscopic depth perception, something that we don't find in visual area 1. With both eyes open these cells are very fussy about the distance to the screen: a given cell responds only if the screen is at the appropriate distance. Evidently the input from both eyes must have a specific alignment for the cell to tire (the principles of stereopsis are explained in Fig. 19). If the screen is not at that distance the cell simply ignores the stimulus and doesn't work at all. There are three categories of cells involved in stereopsis: "near" cells, which respond to stimuli at distances closer than *d;* the distance at which your gaze is fixed; "far" cells, which respond to stimuli at distances greater than *d;* and cells with no disparity. which respond

BRANCHES OF THE VISUAL PATHWAY

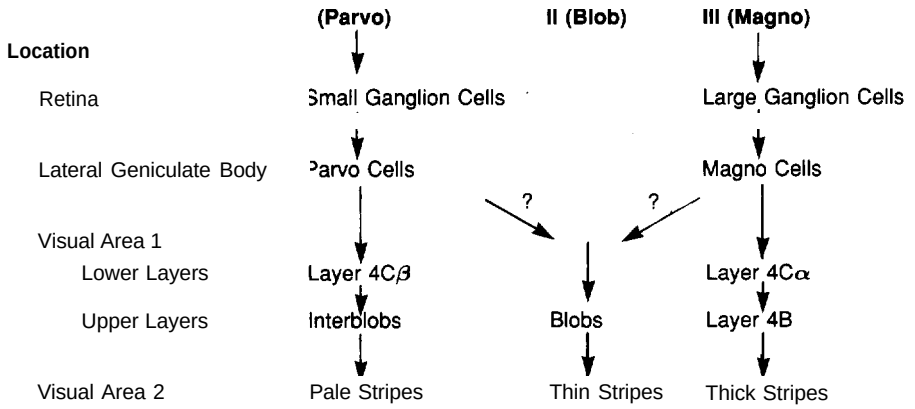

Properties			
Spatial Resolution	**high**	low	low
Orientation Selectivity	**yes**	no	**yes**
Movement Sensitivity	**yes**	no	**yes**
Directionality	no	no	**yes**
Stereopsis	no	no	**yes**
Contrast Sensitivity	low	**high**	**high**
Color	no	**yes**	no

CONCLUSIONS

Pathway 1, the parvocellular, is characterized by high spatial resolution, orientation selectivity, and end-stopping. We guess that this pathway is concerned with high-resolutions form perception.

Pathway II, the blob system, is concerned with color, but not with movement, stereoscopic depth perception, or form.

Pathway Ill, the magnocellular, exhibits systematic selectivity for movement and disparity between inputs from the right and left eye. This pathway thus seems concerned with movement and depth perception.

Fig. 18. Three separate branches in the visual pathway. Resuits of human psychophysical teats support the conclusions above.

to stimuli at the distance *d*. The cells in the thick stripes are thus concerned with stereoscopic depth perception as well as movement.

As you may imagine, the thin stripes, which we know are the terminals of the blobs, are color-coded, and about half of them are involved in the same sort of color problem (color constancy) we described for the blobs. They are thus a continuation of the blob system. The pale regions, finally, are full of end-stopped cells. In visual area 1, you may find that 20 percent of the cells are end-stopped, but in visual area 2 more like 80 percent are end-stopped. So we think that the pale regions are concerned predominantly with form perception. Form perception is a complicated concept, and I am using the term loosely.

All this has many consequences for perception, and I will have time to discuss only two of them before we close. One relates to the fact that the thick stripes are interested in stereoscopic depth perception and get their input from the magnocellular layers. Since, as far as we know, the magnocellular layers are not concerned with color at all, we would predict that stereoscopic depth perception and the perception of movement do not involve color. These predictions can be tested on humans, and I want to show you a few examples of the kinds of tests we have been doing.

First I will show you a simple demonstration to get across the idea that movement doesn't have much to do with color, nor with form. These three parts of vision seem to function independently. You perhaps know that if you turn on a spot on an oscilloscope screen, then turn it off and immediately turn on another spot a small distance away from the first, turn that spot off, and then continue to alternate back and forth, you get a vivid impression of movement from one spot to the other. Psychologists call this apparent movement. This illusion is used in neon signs and in airplane landing strips and of course in cinematography. For some years psychologists have been playing with the array of spots shown in Fig. 20. First you flash two spots on diagonal comers of a square, then turn them off and flash two spots on the other diagonal comers, and so on. Most people looking at this array have the impression that the spots are going up and down, but you can just as well have the impression that the dots

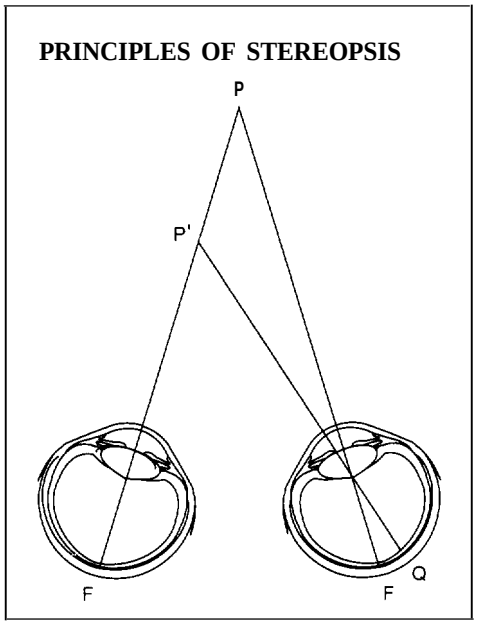

Fig. 19. When an observer fixes his gaze on a point P, the images of P in the two eyes fall on the two foveas. A second point P' closer to the observer than P has Its two images, In this case F and Q', displaced outwards relative to the distance between the two foveas. Simiiarly a point more distant than P wili have its images dispiaced inwards. Relative horizontal displacement of the two retinal images, for near or far objects, is Interpreted by the brain as relative depth. This was discovered in 1838 by Sir Charles Wheatstone, who aiso invented the Wheatstone bridge-and the concertina.

Fig. 20. The film display shows two dots that alternate between diagonally opposite corners of a square. The perception one gets is of dots moving vertically (up and down) or horizontally (from side to side). If the dots are changed to O's and X's, one sees an O going up and down and turning into an X and back to an O or X's and O's going from side to side. It is possible to fiip from one perception to the other. Similarly if the dots are green and red one sees a green dot jump and become a red one **and jump back and become green again. Thus movement perception appears to be distinct from color and form.**

are going from side to side. If you look at the oscilloscope screen long enough, your perception may flip and you will see the spots going up and down rather than horizontally or vice versa. If you have trouble making the flip, say from vertical to horizontal motion, you can block the two bottom dots with your hand, then take your hand away and the dots will seem to be going from side to side rather than up and down. If all of this is so, then you would think that if we made the two top spots X's and the two bottom ones O's you would see them going horizontally; X would go to X and O would go to O. But that's not at all what happens. You are just as happy to see an X going up and down and turning into an O and then back to an X. Similarly if you make two of them green and the other two red, it seems to have no influence on your perception of movement. You are perfectly happy to see a green spot jump over and become red. It is no problem. This seems to suggest that movement perception is quite different than color or form perception.

Now I want to turn to depth perception. I can't show you anything about stereoscopic depth perception without fitting everyone with polarized glasses, but I can show you examples of how other kinds of depth perception, perhaps all types, are colorblind. There are many other cues to depth: occlusion, parallax, movement, and so on. They all seem to fail if the figure you are looking at contains color borders but no change in intensity across these borders. I hope to convince you of that.

To do these demonstrations we use equiluminosity. We take a picture that has red and green areas and try to bal-

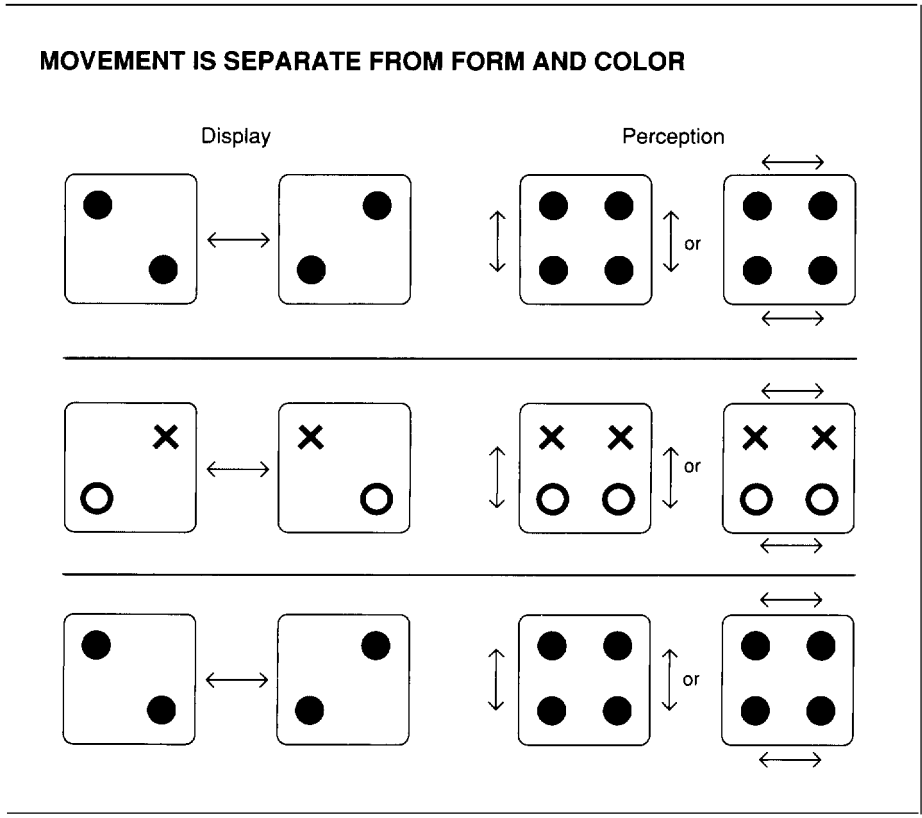

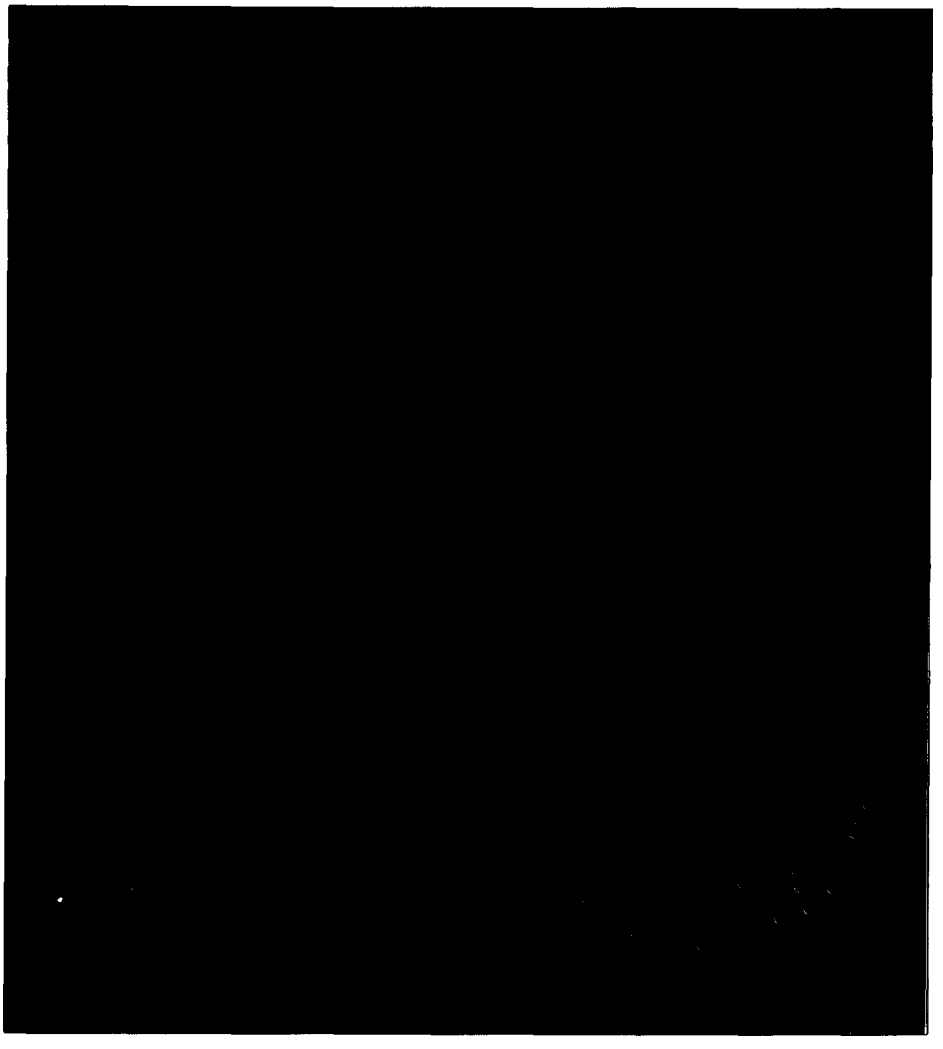

Fig. 21. Dots on a transparent rotating sphere are projected onto a red plane. Two successive images are shown in the figures. When the dots are dark green, the moving dots give the impression of a rotating sphere. When the green dots become equiluminous with the red background the impression of the sphere (that is, of depth) deteriorates, and the dots appear to be moving at random.

ance the two colors so that they are equal in luminance, that is, equally bright, roughly speaking. If you get them equally bright and, say, the system for detecting depth or movement is colorblind, you should stop seeing that phenomenon whereas you would see it perfectly well with a black and white image.

1 have a demonstration in which a bunch of dots are pasted on a rotating transparent sphere, and the moving image is projected on to a red plane (Fig. 21), so the dots on the plane shift in their position in such a way as to imitate the movement of dots pasted on the rotating transparent sphere. Everyone who looks at this film has the impression of a moving sphere. The dots start out dark green, and then gradually become a brighter and brighter green. When you get the brightness of the red background and green dots balanced. your capacity to deduce shape from movement disappears, and all the dots seem to be moving all over the place in a random fashion, You lose the impression of the sphere. Then as I increase the brightness of the green still more, the impression of the rotating sphere returns. This demonstration illustrates why we think that our ability to deduce shape from movement requires luminance differences. When the luminance is equal. color borders alone are *not* enough to give the shape because the visual pathway for color perception is separate from the pathway for perceiving differences in luminance.

In 1984 Cavanagh and Favreau discovered that if you had red and green stripes moving slowly downward on an oscilloscope screen and you change the intensities of the red and green, you can come to some balance for which the impression of movement deteriorates. That is, the red and green stripes seem to be moving more slowly or not moving at all. I want to demonstrate this. What I have here is two sine wave gratings 180 degrees out of phase, one red and one green. We have blocked off half the TV screen with red cellophane, so that on the left half of the screen you see red and black stripes and on the right half you see red and green (Fig. 22). The stripes are moving slowly down the screen. We start with the green stripes lower in intensity than the red and gradually build up the intensity of the green. We have found out that different people have different balance points; that is, the red and green are perceived as equally bright at different relative intensities. This makes the effect hard to demonstrate, and is the reason that I'm slowly building up the green: everyone will get a balance at one or another level of green.

Now if you look at the junction be-

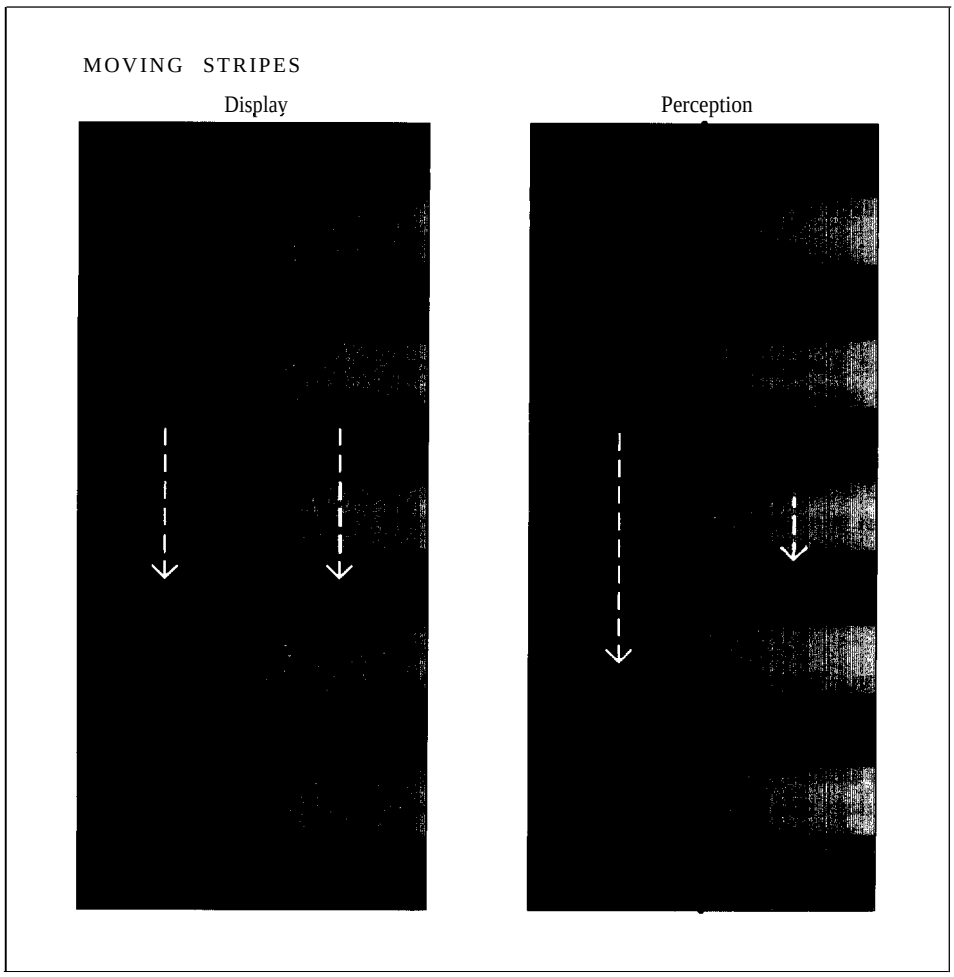

tween the two halves of' the screen, of course you will see that both sets of stripes are moving at the same rate. But if you look away a little bit, say a few feet above the whole demonstration along the middle, and you ask yourself which is moving faster, I think that you will agree that the red and black stripes on the left seem to be moving faster than the red and green stripes on the right. As the relative brightness changes, there may even be some point when the red and green stripes appear to be stationary. But as the film proceeds, we can satisfy ourselves by looking at the border that both sets of' stripes are moving at exactly the same rate. This demonstration clearly suggests that the perception of movement relies heavily on differences in luminance. or brightness. Color cues alone are not reliable.

Fig. 22. Two sets of moving stripes, one red and black, the other red and green, show how movement perception deteriorates when the red and green stripes are equiluminous. Although the stripes move down the screen at equal velocity, when the red and green stripes are equiluminous, they seem to move more slowly than the red and black stripes.

Maybe I have convinced you that perception is more complicated than the word might imply. We get to know better what the words imply by getting a deeper understanding of what is behind them. ■

Questions and Answers

Question: You have focused on a part of' the brain that is transducing external reality inward and allowing us to see it. Have you tried to move into parts of the brain that are doing more abstract things such as problem solving?

Hubel: What I have discussed today represents the kind of investigation that has been done so far, and it is obviously) very far from explaining how you recognize a face, a boat, a hat, or any familiar image. We are very far from understanding what we call shape recognition, to say nothing about more abstract things, such as language, speech, and maybe the most difficult of' all, problem solving. At the moment, the problem of getting to Mars is easy compared to [hat of understanding how the brain solves problems. That is more like the problem of' getting to a planet in the Andromeda galaxy, which is difficult indeed given our life span and the speed of light. on the other hand, understanding perception is not impassible in principle—but we are still a very long way off.

Question: When you look at other parts of the brain, are there any initial clues about how they are organized'?

Hubel: I think there are some, but you might not accept them as bona fide clues. If you look at the organization of the entire cortex, not just visual area 1, you might ask whether the areas responsible for problem solving-the frontal lobes or parietal lobes or something like [ha-are organized differently. To a first approximation they are amazingly similar. Maybe the interesting differences about the organization are still concealed from us because we haven't looked at these areas in the right way, but we simply don't know. The part of the brain that I've talked about comes, roughly speaking, hard-wired: everything I have said today is true in a newborn monkey. Orientation specificity of

individual cells, for example, doesn't take any learning at all. In contrast the areas that are important for languages don't come hard-wired, at least in the sense that we are not born knowing German or any other languages.

Question: Is the auditory part of the brain hard-wired in a manner similar to the visual?
Hubel: We know enough about the auditory system to deduce that inhibition is again going to be important in titrating out excitation to produce stimulus specificity. Although we know a great deal about the response of primary auditory nerve fibers to auditory stimuli, we know very little about the central auditory system, except in a few animals like the bat that use their audition so differently that it may not even be pertinent to our understanding of language. For some reason audition in the central nervous system has been much more difficult to explore, and research has gotten off to a much slower start, but it should prove to be every bit as interesting as vision. From a superficial examination the auditory apparatus in the cortex looks not too dissimilar from the visual cortex. It has an input and an output: it has layers—but this similarity may be just like the similarity of the boxes housing the television set and the personal computer. They look superficially similar, but they are very different, except that both are crummy technologically (I mean the TV and computer!).

Question: Are there physiological differences among people or animals that lead to different responses to the same stimuli?
Hubel: The similarities are more striking than the differences. You have to get into the realm of color and other rather specific things before you find differences between a squirrel monkey and a macaque and a cat. Even though I didn't let on at the time, some of the demonstrations of complex cells that I showed you were actually done in a cat. Only an expert would know that they weren't done in a monkey. The apparatus seems very similar at these early stages in the brain.

Question: Are the cells you have described involved in dreams?
Hubel: My guess is that they are not involved. Dreams are more likely to involve cells that are several or many stages farther into the nervous system, probably in the temporal lobe. The Penfield work, in which the temporal lobe was stimulated and dream-like sequences were produced in epileptic subjects, suggests that this structure is very much involved in vision and dreaming.

Question: What about anesthetics?
Hubel: For the experiments 1 have described, we use a general anesthetic so that the animal is unconscious. But the same things can be tested in waking animals if you implant the electrode and then train the animal to keep its gaze riveted on the screen. In waking, purring cats, for example, or in animals that are walking around and not unhappy, you see no great differences in the response of these cells. General anesthetics probably work primarily on the reticular system, a system deep in the brain, because when you knock out that system either by lesions or concussion or general anesthetic, you lose consciousness. But the anesthetic doesn't have nearly as strong an effect on the visual regions we have tested. The cells do tend to fire more slowly and sluggishly, and we keep checking in chronically prepared animals to be sure that we are not looking at some artifact of the anesthetic. At the moment we have no doubts that the stimulus specificity we have demonstrated is independent of the anesthetic. What is not clear is how far we will be able to penetrate into the nervous system without encountering problems related to the anesthetic. For those experiments we must use chronically prepared animals, which is far more time-consuming and in the end requires more animals.

Question: Are the cells in the visual cortex used for other brain functions?
Hubel: We have not the slightest hint that these cells are involved in auditory stimuli or other sensory inputs except maybe in a very secondary way. If a chronically prepared animal is drowsy and you arouse it, say, with a ringing bell, the cells do respond better but that doesn't mean the response is specifically to an auditory signal. These cells are really very specialized. We have no guarantee that in all cases we have found the optimal stimulus, but after several years of work we are more and more convinced. We have tried numerous things and many of our enemies have too. It is a good thing about enemies; they check up on you.

Question: Do you expect technical advances that might change the rate of progress of your understanding?
Hubel: Technical advances are certainly increasing the rate of progress in anatomy. The methods for revealing the complex, elegant systems of connected structures, such as blobs and stripes have been revolutionized in the last ten or fifteen years. Advances in physiology have come more slowly but I'd really be a pessimist if I thought that big improvements wouldn't come sooner or later.

Question: Do drugs change the sort of picture you've shown us'?
Hubel: We haven't looked very much at the effects of drugs, but other people have shown that certain drugs interfere with the neurotransmitter that is largely responsible for inhibition. In particular, when the drug called bicuculine is

dumped on the cortex, the cells lose their orientation specificity and respond to all orientations. Many groups of people doing pharmacological studies are trying to unravel the visual circuitry by identifying the transmitters responsible for specific responses. This work is still in the very early stages and has yet to solve significant problems, but many people are optimistic about it.

Question: Have you tried to use some complex visual form as a stimulus and compare the responses of someone familiar with this form to someone who is not?

Hubel: We could do something like that by having a person look at something and asking whether the cells fire better when he is paying attention and so on. To do such an experiment with any effectiveness, we would have to be able to record from single human cells without going through the skull, but no technique for that is even in sight. Perhaps twenty years from now we will be able to do such things without requiring an operation. Experiments with waking monkeys have been done, but the animals require a great deal of training. The best work has been done on comparing attentiveness. If the animal is attentive to one part of his visual field and not to another, do the stimuli in the two respective places work differently? They do work differently, but showing that has been a struggle.

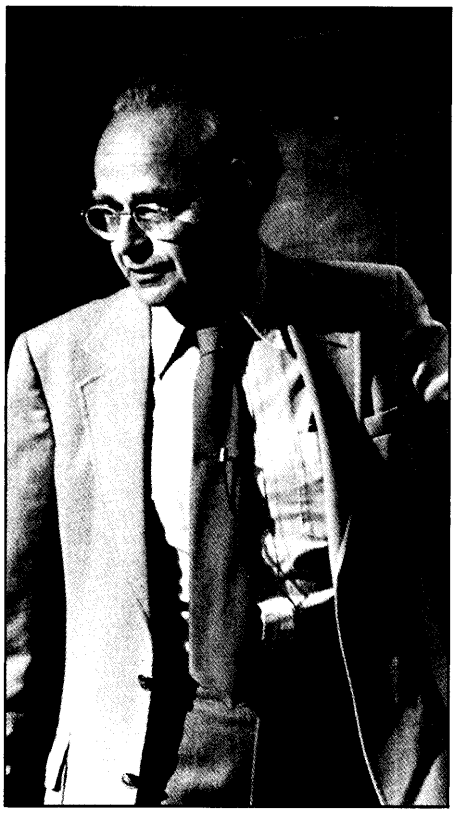

David H. Hubel was born in Canada of American parents, grew up in Montreal, and did honors in mathematics and physics at McGill College. He graduated from McGill Medical School and received training in neurology at Montreal Neurological Institute and Johns Hopkins Hospital. He began his studies of vision in the Neuropsychiatry Division of Walter Reed Army Institute of Research and then returned to Johns Hopkins Hospital to join the laboratory of Stephen Kuffler. There he began a collaboration with Torsten Wiesel that lasted over twenty years and led in 1981 to their receiving the Nobel Prize in Medicine and Physiology. He has been at Harvard Medical School since 1959 and is now the John Enders Professor of Neurobiology. Dr. Hubel descibed himself in *Les Prix Nobel (* Stochholm: Almqvist & Wiksell International, 1982.) as follows: "Since the age of five I have spent a disproportionate amount of time on music, for many years the piano, then recorders, and now the flute. I do woodworking and photography, own a small telescope for astronomy, and I ski and play tennis and Squash. I enjoy learning languages, and have spent untold hours looking up words in French, Japanese and German dictionaries. In the laboratory I enjoy almost everything, including machining, photography, computers, surgery-even neurophysiology ."

Further Reading

David H. Hubel. 1988. *Eye, Brain, and Vision. Scientific* American Library series number 22. New York: Scientific American Library.

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

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

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

We wish to thank Scientific American Library for permitting us to use material from *Eye, Brain, and Vision* in Figs. l–l 1, 14, and 19. Figures 20 and 21 were adapted from the article "Psychophysical evidence for separate channels for the perception of form, color, movement, and depth."