CHAPTER 3

THE EYE AND SEEING

To appreciate how complex the eye is, just talk to an engineer who has tried to build one out of electronic components. The first stage in the visual process—registering the presence of light within an image—is not so difficult. It can be simulated with an array of special transistors called photodiodes. These tiny devices produce voltages that accurately represent the amount of light impinging on them, just like the photoreceptors in your eye. But then what? How do you make any sense out of an array of voltage values? How does the artificial eye figure out what objects in the world produced particular patterns of voltage values in its array of photodiodes?

To accomplish this feat engineers have turned to nature, which holds all the copyrights on the most successful eye designs on earth. Using ideas borrowed from biological eyes, engineers working together with physicians have created wafer-thin silicon microchips whose circuitry mimics some of the neural processing accomplished by the retina (Humayun, de Juan and Dagnelie, 1996). Still, the performance of those artificial retinae pale in comparison to the real thing. At best, the bionic eyes can provide vision capable of supporting limited mobility. In this chapter, we're going to try to think like design engineers, examining the circuitry and neural computations performed by the retina.

By way of preview, here's a partial list of components for the retina of a single human eye: approximately 100 million photoreceptors (rods and cones); 10 million horizontal, amacrine and bipolar cells; and 1.25 million ganglion cells. All these millions of parts are segregated into different layers of the retina, which is no thicker than a postage stamp (recall Figure 2.28). Our goal here is to see how the retina's parts function together to generate biologically useful descriptions of the retinal image for the brain to interpret. To start, let's briefly review what photoreceptors do, since they kick off the series of neural events that culminate in visual perception.

Millions of Points of Light

Each and every photoreceptor performs one simple job: it gauges the amount of light it is receiving. Photoreceptors perform these measurements by generating electrical currents proportional to the amount of light absorbed by the photopigment contained in their outer segments. Over the entire retina, this operation results in millions of separate measurements, each one specifying the light level falling on the tiny region of the eye occupied by a photoreceptor. Viewed this way, we can characterize the photoreceptor matrix as a two-dimensional array consisting of millions of cells. At any given moment, each cell is generating its own unique response to light that expresses the degree to which the electrical current in that photoreceptor cell has been altered. That alteration in current is directly related to the amount of light absorbed by the photoreceptor pigment (Figure 3.1).

But vision doesn't consist of tiny points of light signalled by our photoreceptors. The objects we see in nature aren't like the pictures printed in newspapers, where images are synthesized from clusters of tiny light and dark dots. The array of light measurements registered by the photoreceptors get passed on to a network of neurons that reorganize those millions of raw measurements into more efficient, biologically-relevant messages about the distribution of light in the retinal image. Individual light measurements are transformed into visually important information about the contrast, color, edges and textures, as well as the ingredients, or features, that comprise objects. To understand those transformations, we need to define exactly what is meant by "information about objects."

Think back to our discussion of image formation in the preceding chapter. There we noted that objects absorb some of the light hitting their surfaces and reflect the rest. In effect, objects in our environment "sculpt" the light that eventually arrives at our eyes. The optics of the eye maintain that sculpted spatial structure, forming an image of the sculpted light on the retina. Thus, when it reaches the retina, the amount of light reflected from some object usually differs from the amount of light reflected from the object's surroundings. The key word is *differs*. If there were some way of registering when neighboring retinal regions were being illuminated by *different* amounts of light, one would be on the way to identifying an object's edges or borders—places where the amount of reflected light changes. If one is interested in objects and edges, one is *not* interested in regions of the retina over which the light level remains constant. Homogeneous, uniformly illuminated regions of the retina probably do *not* represent the image of an edge. To identify an edge, the retina needs to note where there are *differences* between the light levels at adjacent locations. As you'll see, many retinal cells are designed to do precisely that: respond to differences between adjacent levels of light.

Our immediate goal, then, is to understand how the retina condenses and reorganizes the multitude of messages supplied by the photoreceptors. Then we will consider how the condensation and reorganization actually affect the appearance of objects and their surfaces.

Eventually, we need to examine the neurons forming the middle layers of the retina: the bipolar, amacrine and horizontal cells. They are the circuit elements that radically transform the photoreceptors' signals into something much different from light measurements. But for the moment, let's skip to the **retinal ganglion cells,** the neurons responsible for the last stage of processing within the eye itself. Once you understand what the ganglion cells do, you'll find it easier to see how those middle layer cells make the visual process possible.

During our discussion, keep in mind two important facts about the ganglion cells. First, although they do respond to visual stimulation, they do not themselves absorb light; they are *not* photoreceptors. Ganglion cells process neural information that the other retinal cells have received from the photoreceptors. Without input from

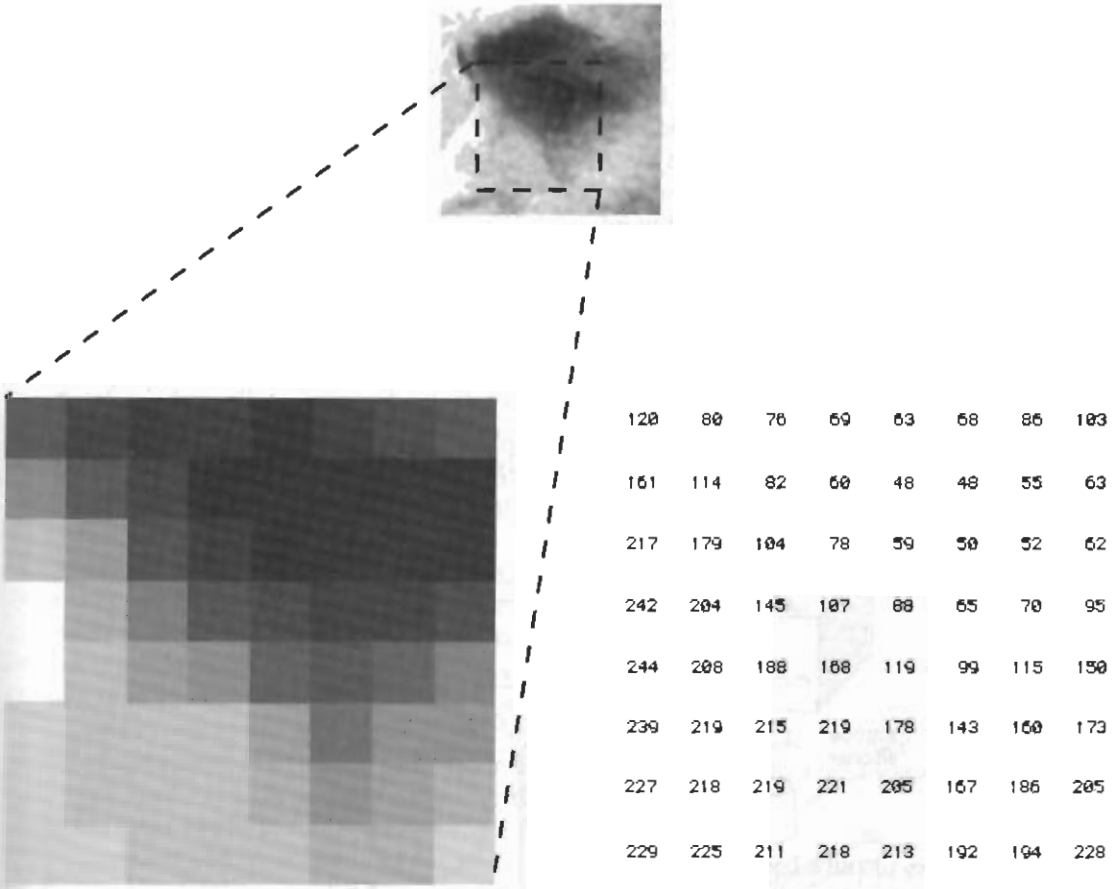

FIGURE 3.1
Photoreceptors generate electrical signals proportional to the amount of light impinging on them. The array of spots on the left depicts a portion of the image of a human face as "seen" by an array of photoreceptors. The matrix on the right contains numerical values representing electrical signals. The value of each "signal" is proportional to the intensity of light striking a given photoreceptor. Notice the spatial correspondence between intensity levels and matrix signal values.

those other cells, the ganglion cells would be blind to everything happening in the visual world.

Second, ganglion cells can only signal the outcome of their processing by generating **action potentials,** brief electrical discharges carried by the nerve fibers of the ganglion cells to the central visual stages within the brain. So whatever a ganglion cell has to "say" about a visual stimulus must be expressed using this one-"word" vocabulary. This restriction actually applies to *all* further stages of visual processing—neurons talk to one another in a language composed entirely of action potentials, or neural impulses, as they are sometimes called. Particular neurons speak with a burst of impulses only when certain types of visual stimuli appear. By virtue of their early position in this chain of visual processing, retinal ganglion cells set this neural dialogue in motion. Let's examine now what the ganglion cells have to say.

THE RETINAL GANGLION CELLS

N Neural

The human eye contains an estimated 1.25 million retinal ganglion cells. Comparing this figure to the roughly 100 million receptors in the eye, you know from the outset that ganglion cells must be coudensing the raw messages from the receptors. Imagine you are handed a 1,000-word essay (about 4 double-spaced, typed pages) and told to edit it down to 12 or 13 words (1 or 2 sentences)—without losing the essentials of its message. To meet this challenge, you must identify the essay's major points and then rephrase them, condensing in a way that preserves the essence of the original. Retinal ganglion cells face the same kind of problem: they must collate messages from the more numerous photoreceptors and—with help from the horizontal, amacrine and bipolar cells—summarize those messages in a biologically relevant way. How do the ganglion cells accomplish this?

E Execute

The most direct way to tackle this question is to determine what kinds of visual stimuli are best able to activate ganglion cells. Figure 3.2 illustrates the experimental procedure for determining a visual cell's preferred stimulus. An experimental animal, a monkey in this case, is shown facing a screen. A tiny, fine-tipped wire called a **microelectrode** is surgically and painlessly inserted into the part of the visual system under study, in this case, the optic nerve (which com-

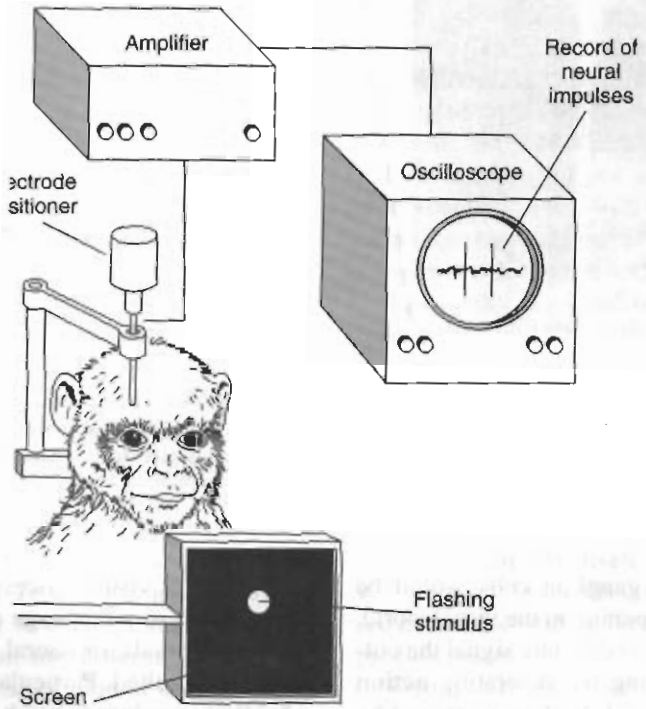

FIGURE 3.2
Laboratory setup for recording action potentials from single neurons. The placement of the recording electrode governs which stage of the visual nervous system will be examined. In the case shown, the electrode would pick up action potentials generated in the axons of the retinal ganglion cells.

prises the axons of the ganglion cells). The probe can be positioned close enough to an individual ganglion cell **axon** (the portion of the cell that carries its impulses out of the retina) so that the electrode picks up the action potentials (neural impulses) arising from just that cell. One can then monitor the number of action potentials generated by this single cell, and try to influence the cell's activity level by presenting various sorts of visual stimuli on the screen. This technique, called **single cell recording,** has been successfully employed to determine what kinds of visual stimuli activate cells at different stages within the visual system. Here we are interested in the retina. In effect, we wish to ask the ganglion cells "What sort of visual stimulus do you reliably respond to?"

Before presenting *anything* on the screen, you discover that the ganglion cell is already active. The electrode is picking up an irregular but persistent chatter of action potentials from the cell. This spontaneous activity continues even when the monkey is in complete darkness. The occurrence of individual action potentials over time is plotted in Figure 3.3. In each of the three panels, small vertical lines represent single action potentials from one retinal ganglion cell. Time is traced out along the horizontal axis. Looking at panel A, you see the impulses occurring in the absence of visual stimulation—this is the cell's spontaneous activity.

Since it is spontaneously active when no light is present, the cell must signal the presence of light by a change in its *level* of activity. Suppose your job is to discover what it takes to produce that *change* from spontaneous activity. Knowing that photoreceptors (from which the ganglion cells receive input via those middle layer cells) are small, you start by moving a small spot of light around over the screen in front of the monkey. By doing this, you are moving the light over the monkey's retina, stimulating photoreceptors wherever you move the spot. So, in effect, your job is to search for an area of the retina where the image of the spot of light will influence the ganglion cell's level of activity.

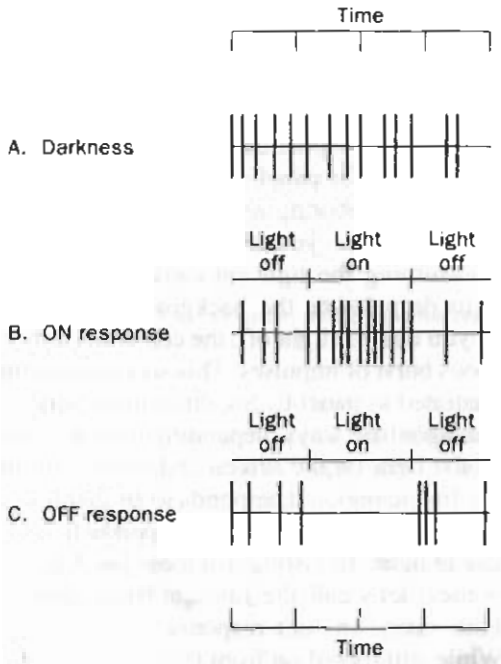

FIGURE 3.3
Neural activity (shown as vertical lines) of retinal ganglion cells. See the text for a full discussion.

As we proceed with our example, it is very important to keep the following in mind: different regions on the screen in front of the monkey correspond to different areas on the monkey's retina. To maintain this correspondence, the monkey's eye must remain perfectly still. If the monkey moved its eye around during the experiment, you could never be certain when your spot of light fell on its retina. Monkeys can be trained to hold their fixation very steady, but sometimes investigators use a drug that temporarily stabilizes the eye during the experiment. Once the monkey's eye is immobilized, positioning the spot in the center of the screen places the image of the spot on the center of the monkey's eye, in the macula. With this in mind, let's start our experiment.

By moving the spot of light around on the screen, you discover a region of the retina where the spot causes an increase in the recording cell's

activity. Concentrating on this region, you find that the cell generates a burst of impulses when you turn the light on this area. When you turn the light off, the cell's activity quickly settles back to its background (spontaneous) level. This outcome is shown in panel B of Figure 3.3. Next, you test a neighboring area of the retina in the same way. Now you find just the opposite result—turning the light on causes the activity level to drop *below* the background level. But when you turn the light off, the cell emits a short, vigorous burst of impulses. This second outcome is illustrated in panel C. So, this cell responds in two antagonistic ways, depending on where you place the light on the screen and, hence, on the retina. In one region it responds to an *increase* in light, whereas in the other it responds to a *decrease* in light. To distinguish these two kinds of responses, let's call the first an "ON response" and the second an "OFF response."

While still recording from this same cell, suppose you now test at other, nearby locations on the retina. You find that ON responses can be elicited from anywhere within a restricted, circular region. In fact, enlarging your spot so it just fills this circular region produces a very vigorous response. The regions giving an OFF response, however, form a ring that completely surrounds the circular ON region. So the same spot of light has opposite, or antagonistic, effects in the center versus the surrounding area, usually shortened to the surround. If you label these two regions using plus signs (for ON) and minus signs (for OFF), the composite looks like a small circle of plus signs surrounded by a ring of minus signs, as shown in panel A of Figure 3.4.

Light placed anywhere *outside* this donut-shaped composite has no influence whatsoever on this cell's activity. In other words, only light falling within this restricted, concentrically shaped area of the retina is registered by this ganglion cell. This area constitutes that cell's **receptive field**—the patch of retina within which a cell's activity may be influenced. (The term "receptive field" was coined by H. K. Hartline [1938] in his classic work on retinal ganglion cells in frogs.) The concept of a receptive field is extremely important for understanding visual processing. As you will come to appreciate in this chapter and the next one, the receptive field serves as a kind of template with which a cell gauges the pattern of light falling within a restricted area of the retina. That pattern of light must fit into the receptive field in order to activate a cell.

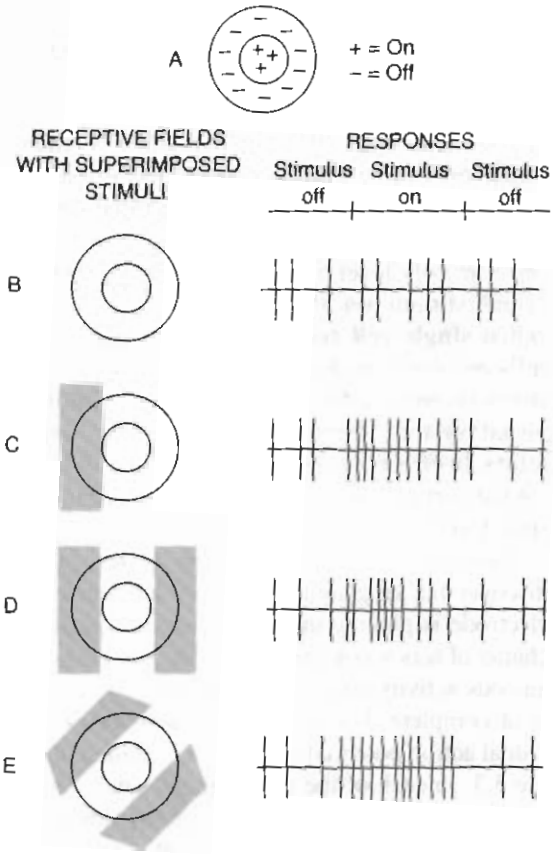

FIGURE 3.4
A single ON-center retinal ganglion cell (A) responding to uniform illumination (B), to a dark/light edge (C), to a vertical bar of light (D), and to an oblique bar of light (E).

To illustrate what we mean, consider the receptive field mapped out in panel A of Figure 3.4. Can you picture an optimal visual stimulus for this cell, one that would produce the most vigorous increase in neural activity? First, imagine illuminating the entire retina, in effect filling the cell's receptive field with light. As shown in panel B of Figure 3.4, the cell gives only a weak response to uniform illumination. This is because such a stimulus produces opposite effects in the center and the surrounding area. The two antagonistic regions compete with one another, resulting in a near stand-off. This interaction between antagonistic regions is called **lateral inhibition.**

Now imagine what happens when an edge is positioned in the manner shown in panel C. The ON-center portion of the receptive field receives an increase in light, its preferred stimulus, while a good portion of the surround receives a reduced level of light, its preferred stimulus. The net result is a vigorous response from the cell. As panels D and E show, the cell would also respond well to bars of light positioned appropriately within the receptive field. Incidentally, because these center/surround areas are nearly always concentrically arranged, ganglion cells will respond well regardless of whether the edge is oriented vertically, horizontally, or diagonally. The orientation of an edge or bar is irrelevant so long as the edge or bar is positioned appropriately within the receptive field.

This antagonistic arrangement of center and surround within the receptive field enables the retinal ganglion cell to perform its filtering work. Lateral inhibition has enabled the cell to condense the messages from a patch of photoreceptors into a single statement: "I detect a light/dark boundary." By accenting the *difference* in light levels on adjacent areas of the retina, the cell has begun the process of extracting perceptually relevant information.

Nearly all ganglion cells have concentrically arrauged receptive fields, composed of a center and a surround that respond in an antagonistic

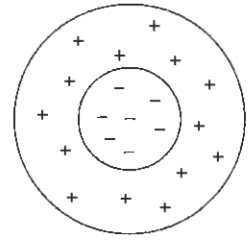

FIGURE 3.5
A receptive field of an OFF-center retinal ganglion cell.

fashion. Some of those cells have ON centers and OFF surrounds, like the receptive field illustrated in Figure 3.4. (Box 3.1 discusses one interesting perceptual phenomenon thought to result from these cells.) Other cells have just the opposite layout, with an OFF center and an ON surround. An example of this latter type of cell is shown in Figure 3.5.

There are about as many ON-center cells as there are OFF-center cells. Both types respond best to light/dark boundaries. Together, the 1.25 million receptive fields of the ON-center cells and OFF-center cells are interwoven to comprise a mosaic covering the entire retina. As a result, whenever light falls on a limited patch of the retina, that light is bound to affect a number of retinal ganglion cells of both types, producing opposite effects in the two. Don't imagine, though, that these opposite effects cancel one another. At higher stages of the visual system, information from ON-center and OFF-center cells remains segregated, allowing information from both types to be used (Wässle, Peichl, and Boycott, 1981). Moreover, behavioral studies of animals in which activity within the ON cells has been chemically disrupted reveal that the ON and OFF cells support different aspects of vision (Schiller, Sandell and Maunsell, 1986). With the ON cells inactivated, animals have difficulty detecting increments in light, whereas light decrements are readily perceived. Low-contrast objects are also more difficult to see without a functioning ON system.

BOX 3.1

BLACKER THAN BLACK

Because of its spontaneous activity, an ON-center cell sends the brain a stronger message when no light falls in its receptive field than it does when its surround alone is illuminated. This curious state of affairs suggests the possibility that some light may actually appear darker than no light at all. In fact, surrounding a dim area with a sufficiently intense light makes that dim area appear darker than an area that contains no light whatever. Probably the intense surround drives the activity of ON-center ganglion cells *below* their spontaneous levels, signaling the brain that something blacker than black is present (Brown and Mueller, 1965).

The consequences of this spontaneous activity show up in everyday life. Think about what it's like to wake up in the middle of the night in an absolutely dark room. Usually the room doesn't appear totally black. In fact, many people experience dim, illusory, swirling lights, the result of spontaneous activity in the visual system (Hurvich and Jameson, 1966).

Some artists exaggerate discontinuities in intensity in order to highlight the outline of figures in their work. If an artist wants to create the deepest possible black region in some painting, he or she must do more thau simply use black paint. Even if that black paint reflected no light at all (which isn't really possible with paint), receptive fields in which the black paint was imaged would still be sending spontaneous messages to the brain. To reduce those messages to a minimum, the artist surrounds the black paint with an area of white or other light-colored paint. The contrast between the two areas intensifies the blackness produced by the dark paint. Ratliff (1972) contains a good introduction to the uses of lightness illusions in art.

The same effect is exploited by your television set and your computer monitor. The screen cannot get any darker than it is when the power is off. Yet the screen doesn't look black when the video monitor is off—it looks dark grey. So how is it that black objects are produced when the monitor is on? Gray regions appear black when surrounded by other, light-colored regions.

Some parents and children also take advantage of light's ability to create darkness, often unwittingly. Many children have a hard time falling asleep unless conditions are just right. Not only must it be past their appointed bedtime, but conditions outside must also confirm that it is bedtime—it must *look* sufficiently dark outside. Some wise parents take advantage of a lightness illusion to hasten bedtime—as some children realize. The Finnish-American poet Anselm Hollo captured this idea by putting the following words into the mouth of a 4-year-old:

switch on the light
so it gets dark outside
and we can go
to bed.
(1977, p. 30)

The Neural Architecture of ON and OFF Regions

We have more to learn about ganglion cells and their center/surround properties. But now is a good time to look more closely at the cells in the intermediate layers of the retina, for it is those cells that are responsible for the spatial layout of ON and OFF regions. Look at Figure 3.6, a highly magnified schematic of the arrangement of the retina's parts at one small location. Let's concentrate on the three types of cells that lay between the photoreceptors and the ganglion cells.

First, notice that the photoreceptors (shown pointing upward in this drawing) are interconnected by laterally spreading cells, aptly named "horizontal cells." Any horizontal cell may be in contact with dozens of neighboring photoreceptors, and each photoreceptor makes contact with several horizontal cells. Each and every horizon-

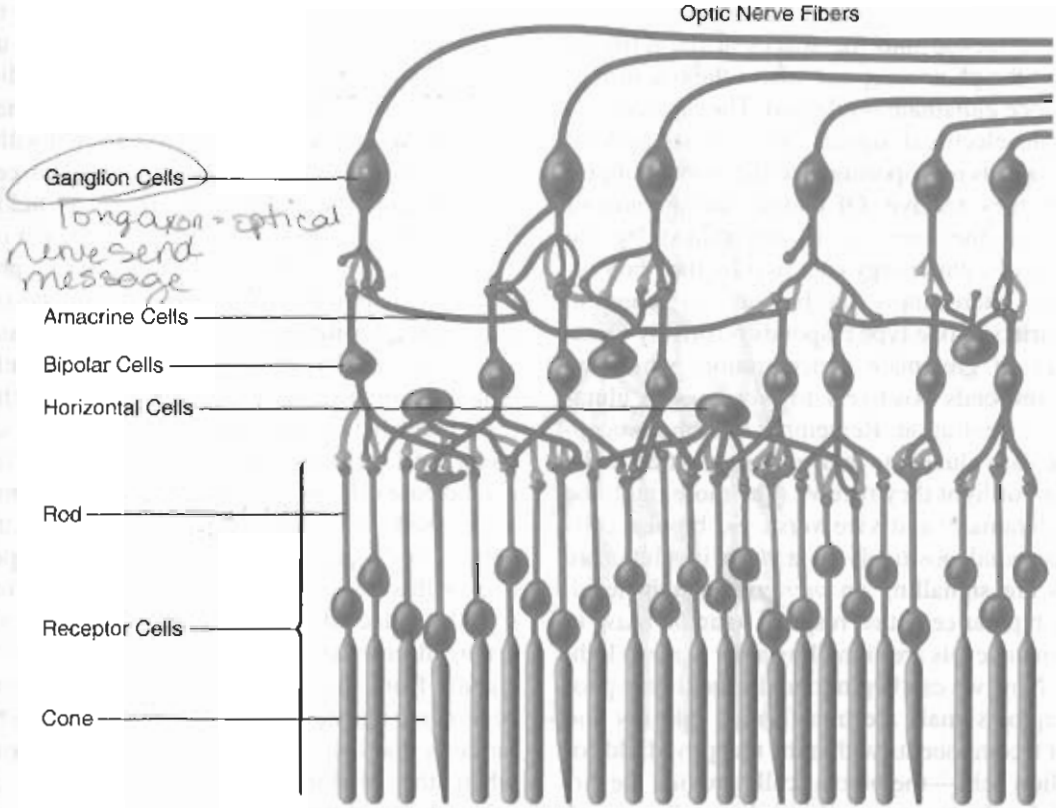

FIGURE 3.6
A magnified view of retinal cells.

tal cell does only one thing: it modifies the strength of the signals generated by its neighboring photoreceptors.

It's an ingenious arrangement. Any one horizontal cell receives input from a cluster of photoreceptors and, in turn, provides input back to those very photoreceptors. Horizontal cells turn down (attenuate) the strength of the signals generated by those individual photoreceptors. This attenuation most strongly affects those photoreceptors whose signals are the weakest to begin with. The more active photoreceptors are less affected and, hence, their signals become *relatively* larger. Thus, when one small cluster of photoreceptors within the neighborhood is more active than the rest, the horizontal cell accentuates the strength of their signals relative to the signals in the other, less active photoreceptors. This feedback arrangement accentuates differences in photoreceptor signals, differences that can be traced to locations in the retinal image where there are transitions in the spatial distribution of light. Horizontal cells, in other words, amplify photoreceptor signals associated with edges—they are a fundamental part of the lateral inhibition described above.

But how are these modified photoreceptor signals conveyed to the ganglion cells? Here's where the bipolar cells enter the picture. As you can see in Figure 3.6, they, too, make synaptic

contact with the photoreceptors. Their receiving ends are tucked into the spaces at the terminal ends of the photoreceptors where the transmitter substance glutamate is released. The bipolar cells generate electrical signals, and the strength of those signals is proportional to the amount of glutamate they receive. Of course, the amount of glutamate they receive is determined by the amount of light energy captured by the photoreceptors. Interestingly, the bipolar cells come in two varieties: one type responds positively to *decreases* in glutamate concentration, while the other responds positively to *increases* in glutamate concentration. Remember that photoreceptors release glutamate in inverse proportion to the amount of light they receive. The more light, the less glutamate, and vice versa. So bipolar cells that respond positively to a *drop* in glutamate levels are signalling an *increase* in light level, while bipolar cells that respond to an *increase* in glutamate levels are signalling a *decrease* in light level. Now we can begin to understand how photoreceptor signals are transformed into ON and the OFF components within the receptive fields of ganglion cells—the bipolar cells provide the circuitry for the two types of ganglion cells.

Actually, there are at least a dozen different kinds of bipolar cells in the human retina. They differ in both size and the patterns of connections they make with rod and cone photoreceptors and with ganglion cells. But they all do essentially the same thing: recombine photoreceptor signals (which have been modified by horizontal cell connections) and pass those recombined signals to the ganglion cells.

What about the amacrine cells? Relatively little is known about their role in retinal processing. We do know that amacrine cells receive inputs from bipolar cells and, in turn, modify the responses of those bipolar cells. (In this respect, amacrine cells behave like feedback cells, comparable to the horizontal cells.) It is thought that amacrine cells may influence the temporal dynamics of ganglion cells, meaning how vigorously they respond over time.

In summary, the network of cells in the retina's intermediate layer form the circuitry underlying the center/surround layout of ganglion cell receptive fields. Figure 3.7 shows how these elements are thought to be interconnected within the circuitry for one, ON-center ganglion cell. The central portion of the cell's receptive field is "constructed" from signals collected from a relatively small, circular patch of neighboring photoreceptors. For an ON-center cell like the one we're considering, those photoreceptors' signals are registered by inverting-type bipolar cells, which become more electrically positive when their contributing photoreceptors become more electrically negative (as they do when light levels increase). These positive electrical responses are passed on to the ganglion cell, giving that cell its ON-center response. The surround portion of this ON-cell's receptive field arises from signals collected from a larger, concentrically arranged circular patch of photoreceptors. The signals from this larger collection of photoreceptors are registered by noninverting-type bipolar cells that become electrically more negative when the photoreceptors decrease their responses. These negative electrical responses are passed on to the ganglion cell, giving the cell its OFF-surround response. Note, by the way, that the central interior of this surrounding region includes the same photoreceptors that contribute to the center.

To show the net outcome of this center/surround arrangement, we can plot the sensitivity of the ON-center and OFF-surrounding regions in the manner shown in Figure 3.8. It is conventional to use a pair of bell-shaped curves (called Gaussian curves) to depict the responsiveness of the center and surrounding portions. The center and surrounding Gaussian curves subtract from one another, yielding a resultant called the difference of Gaussians (or **DoG curve,** for short). A DoG curve provides a good approximation of the spatial layout of most retinal ganglion cells. It is simple to apply the same scheme to the construction of an OFF-center receptive field, just by inverting

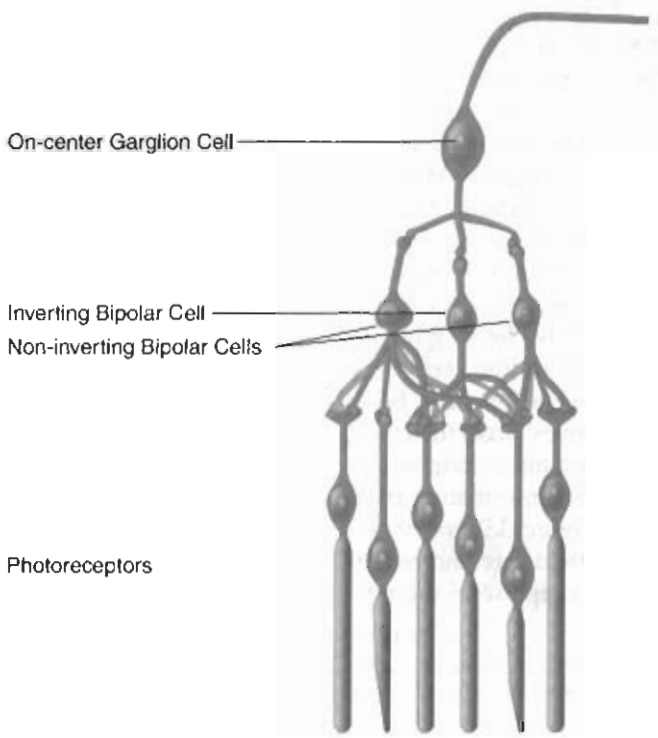

FIGURE 3.7
Schematic showing how a center/surrounding area receptive field is "constructed" through interconnecting retinal cells.

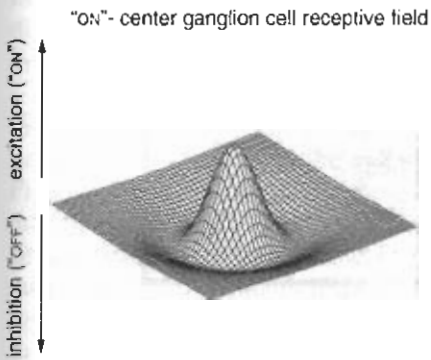

FIGURE 3.8
Sensitivity profile of a retinal ganglion cell with center/surrounding area organization.

the signs of the two Gaussians contributing to the center and surround.

This completes our overview of the intermediate circuitry within the retina, and we'll be referring back to this stage of processing in subsequent discussions. But armed with this expanded view, we're now ready to return to the ganglion cells, to learn more about the neural messages they encode and transmit to the brain. To do this, let's imagine that we've moved our microelectrode from ganglion cell to ganglion cell, determining what visual pattern of light produces the most vigorous responses in each cell. In the course of this exercise, we would discover several important receptive field properties.

Receptive Field Size

The sizes of receptive fields vary systematically with retinal location (Wiesel and Hubel, 1960; de Monasterio and Gouras, 1975). Receptive fields in the center of the retina, within the macula, are quite small, with some on the order of 0.01 millimeters. Cells with these centrally placed receptive fields monitor tiny areas of the visual world wherever the monkey is looking. As you move away from the macula into the periphery of the retina, you find that the receptive fields grow increasingly larger. In fact, 10 millimeters away from the fovea, receptive field centers may be as large as 0.5 millimeters, 50 times larger than their foveal counterparts. Cells with these peripherally placed receptive fields collect information from larger areas of the retina. Figure 3.9 graphs the relation between receptive field size and retinal location. The horizontal axis represents the retinal location around which the receptive field was centered; "fovea" stands for the center of the macula. The vertical axis is the diameter of the center portion of the receptive field. Note that with increasing **eccentricity**—deviation from the center of the retina—the size of receptive fields tends to increase.

In and around the fovea, receptive fields are quite small, so it naturally takes more ganglion cells to cover a given area of retina in the fovea compared to more peripheral parts of the retina. In fact, the fovea accounts for only about 2 percent of the total retinal area, whereas about 33 percent of all ganglion cells are devoted to the neural analysis of image features within this small, but critically important, retinal territory.

The graph in Figure 3.9 also reveals another important fact. At each of the various eccentricities, we have data for more than just a single re-

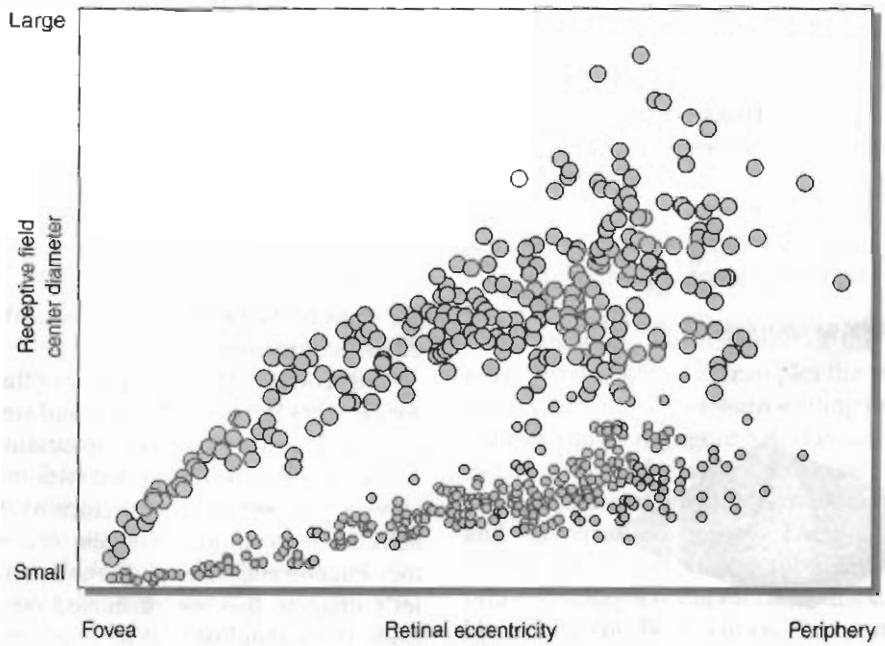

FIGURE 3.9
Graph showing how the sizes of receptive field centers increase with distance from the fovea. The size of the entire receptive field is larger than just the center diameter, since the entire receptive field consists of a center *and* a surround.

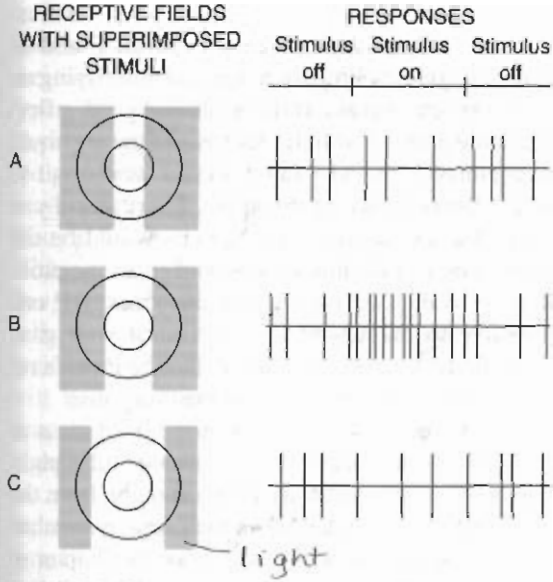

FIGURE 3.10
A single ON-center receptive field responding to bars of light of varying width.

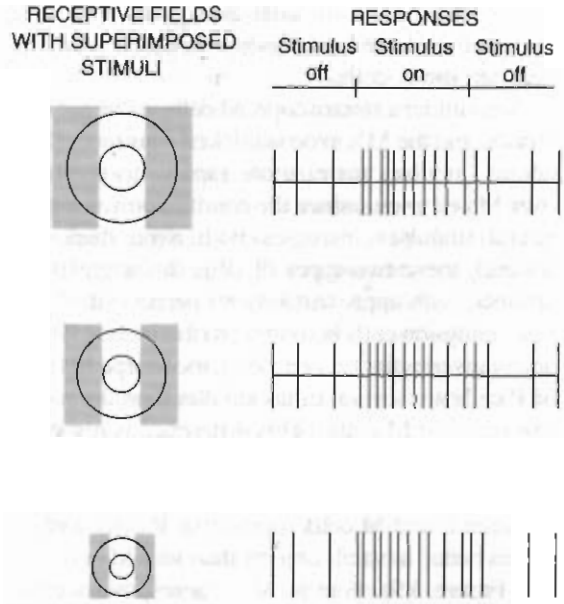

FIGURE 3.11
Three ON-center receptive fields with different-sized receptive fields responding to bars of light of matching width.

ceptive field. Comparing data collected at the same eccentricity, we see that not all receptive fields at that eccentricity have exactly the same size. Because these size differences occur at a single locality on the retina, the phenomenon is known as *local variation.* So the sizes of receptive fields differ in two ways: first, they vary with retinal eccentricity; second, they vary locally. In fact, we see two clusters of receptive field sizes at different eccentricities, a point we shall return to in a moment.

Think what this variation in receptive field size means for the sort of stimulus that would be best suited to activate a particular cell. As shown in Figure 3.10, there is one bar width size (panel B) that elicits the best response from that cell. Bars smaller (panel A) or larger (panel C) than this produce a less than optimum response. As you can see in Figure 3.11, cells with small receptive fields will respond best to small objects, while those with large receptive fields will prefer larger objects. This principle suggests that analysis of object size may be inaugurated in the retina.

Three Types of Retinal Ganglion Cells

So far we've distinguished retinal ganglion cells on the basis of receptive field size and center type (ON versus OFF). But retinal ganglion cells differ in other ways, some of which are particularly relevant for understanding vision. These differences form the bases for classifying ganglion cells into one of three groups, the **M cells,** the **P cells** and the **K cells** (Hendry and Calkins, 1998). *M* stands for magnocellular, *P* for parvocellular* and *K* for koniocellular. These three terms refer to the three classes of brain cells to which the M, P and K cells relay their ueural impulses. Let's survey some major differences among these three cell

*Some researchers refer to the P and M ganglion cells as "parasol" and "midget" cells, respectively. These terms are derived from the anatomical shapes of the two classes of ganglion cells.

types. We'll start by contrasting the M and P cells, which have been studied in much more detail than the K cells.

Seen under a microscope, M cells are larger than P cells, and the M's axon is thicker. This means that neural impulses travel more rapidly to the brain over M cell axons, since the conduction velocity of neural impulses increases with axon thickness. Second, these two types of cells differ greatly in number, with approximately 80 percent of all primate ganglion cells belonging to the P class. Third, at any given retinal eccentricity, the receptive fields of P cells are several times smaller than the receptive fields of M cells. (This difference is responsible for much of the vertical scatter in receptive field size seen in Figure 3.9.) The difference in field size between P and M cells means that P cells will respond better to small objects than will M cells (recall Figure 3.9). Fourth, M cells respond well to very small differences in light levels in center and surround. P cells require greater contrasting difference in light between center and surround portions of the receptive field before they will respond strongly (Kaplan, Shapley, and Purpura, 1988). This distinction between cell types suggests that M cells may be especially important for the perception of objects of low contrast, such as dark gray letters on a medium gray background. P cells may be more important for seeing high-contrast objects, such as black letters on a white background. Fifth, M cells respond well even when a visual stimulus is turned on and off very quickly, whereas P cells respond poorly if at all to such a temporally modulated stimulus. This difference in temporal sensitivity means that M cells are better able to register transient visual events, such as the presence of a rapidly moving object.

The sixth difference between M and P cells is especially intriguing, since it probably relates to color vision. For P cells, an excitatory response is evoked only when the receptive field is stimulated with light of a particular color (for example, red). P cells are inhibited by the presence of another, quite different color (for example, green). For M cells, in comparison, color makes no difference. They respond to light regardless of color. To illustrate this important difference between P and M cells, imagine looking for a ripe, red apple lying on a lawn of green grass. If the apple and grass reflect the same amount of light (that is, if they are equal in lightness), the M ganglion cells would fail to signal the presence of the apple. Everywhere you look, the responses of the M cells would be the same. The P cells, however, would solve the problem: as your gaze roams over the grass, a P cell sensitive to red would be silent until your gaze brought the image of the apple into the P cell's receptive field, thus evoking a robust response.

Not as much is known about the third class of cells, the ones comprising the koniocellular pathway. Cells in this pathway get their name from the neurons that they innervate in the lateral geniculate nucleus, which we will learn about in Chapter 4. Those geniculate neurons are extremely small (*konis* is Greek for "dust") and, therefore, very difficult to isolate for physiological study. In terms of their response properties, K cells more closely resemble P cells, in that their responses depend on the color of light within their receptive fields. In particular, most K cells are excited by blue light and inhibited by yellow light (Dobkins, 2000). K cells may also play a role in shutting down vision temporarily whenever we move or blink our eyes. Detailed reviews of the K cells are provided by Casagrande (1994) and by Hendry and Reid (2000). Table 1 summarizes these three classes of cells.

TABLE 1
Physiological Properties of P, M and K Cells

Characteristic	P cells	M cells	K cells
Cell size	Small	Large	Very small
Conduction velocity	Slow	Fast	Slow/variable
Cell population percentage	80%	10%	>10%
Spatial resolution	High	Low	Moderate
Temporal resolution	Low	High	High
Contrast sensitivity	Low	Good	Modest

Although P, M and K cells constitute the lion's share of neurons in the retina (over 90 percent of the total), there are other, less researched cell types. These other cell types send their axons to phylogenetically older regions of the brain that are probably involved in the control of eye and head movements. These other cell types are also thought to provide the visual information that drives circadian rhythms and influences your sleep-wake cycle. In contrast, the P, M and K cells, project to higher, phylogenetically newer brain centers mediating visual perception. As you will learn, nature has gone to great lengths to keep separate the information carried to these higher centers by these three classes of cells. In the next chapter, we will learn more about the separate visual pathways that originate from the retina's P, M and K cells, and will discuss further their possible roles in visual perception.

This completes, then, our abbreviated survey of retinal ganglion cells. By examining the workings of the retinal ganglion cells, you have learned several important things about the eye's processing of visual information. You now know that this processing begins with the photoreceptors, which act like an array of tiny photocells, each specifying the level of light falling within the purview of the photoreceptor. These 100 million messages about light intensity are then passed on to a complex network of intermediary cells that integrate information from groups of neighboring photoreceptors. The results of this integration are conveyed to the retinal ganglion cells. Because of the center/surround organization of their receptive fields, the vast majority of ganglion cells are designed to detect differences in light level, or *contrast*, as it is called. These cells are much less concerned with the overall level of light.

In some of these cells, this center/surround organization is designed to extract information about color contrast. This kind of local receptive field analysis is performed over the entire retina by the 1.25 million or so ganglion cells in the eye. Hence, everything you see must have registered its presence within this retinal machinery. The particulars of this machinery necessarily influence the *way* you see. There is no other route to visual perception but through the retinal ganglion cells.

The remainder of this chapter relates certain properties of vision to events occurring in the retina. A few of these events transpire in the photoreceptors, while others occur at the level of the retinal ganglion cells. It's important that you understand what we mean when we claim that some property of vision is caused by the idiosyncrasies of some retinal cell. We are *not* saying that conscious visual perception occurs in the retina. Most visual scientists think that the processes underlying what we call "vision" actually take place in the brain not in the eye.

Instead, we're saying that events in the retina *shape* vision by emphasizing some information (such as differences in light level) and by deemphasizing other information (such as uniformities). For example, retinal ganglion cells respond very strongly to discontinuities in illumination, and this operating principle reveals itself, sometimes vividly, in visual experience. Potentially important information, in other words, gets accentuated. In contrast, retinal cells fail to respond to images composed of wavelengths longer than 700 nanometers, meaning that the brain receives no information about such images. So some information gets ignored. Thus, even though sight occurs in the brain, the retina preordains much of what we can and cannot see. Brindley (1970) and Teller (1989) discuss the logical bases for attributing perceptual events to the behavior of certain physiological processes.

PERCEPTUAL CONSEQUENCES OF CENTER/SURROUND ANTAGONISM

The preceding section emphasized the antagonism between the center and the surround of a retinal ganglion cell's receptive field. The net response of such a cell is the sum of these two opposing influences. This antagonistic arrangement reorganizes the receptors' raw information into neural signals

associated with light/dark boundaries (such as with edges and contours). In the process, center/surround antagonism accentuates those boundaries and, at the same time, de-emphasizes regions where light intensity is uniform. This differential boost given to edge information has led to the popular but controversial idea that the perceived lightness over regions of a surface is related to activity in retinal ganglion cells. In the following sections, we consider a couple of intriguing perceptual illusions that explore the operation of center/surround antagonism in human vision.

From the outset of this discussion, we need to clarify the distinction between the terms "brightness" and "lightness." They're often used interchangeably, but they should be differentiated. Among vision experts, the term "brightness" refers to the amount of light that appears to be arising from a given spatial location. For instance, we may speak of the brightness of the light reflected from, say, a piece of white paper. Brightness also refers to the perceived level of illumination produced by an emitting source. For instance, we may speak of the brightness of a room light or a star. "Lightness" is a property of object surfaces illuminated by light; it depends on the amount of light reflected from the surface of an object (which itself depends on both the amount of light illuminating the object and the amount of that illumination reflected by the surface). Lightness ranges from black, through shades of gray to white. A sheet of white paper is always lighter than a sheet of gray paper, but the gray sheet may look brighter when it is in direct sunlight compared to being in the shade. To reiterate, brightness is a property of light (reflected or emitted), while lightness is a property of surfaces. In most instances, it is the lightness of a surface that we are concerned with. Our everyday activities entail recognizing objects, not light, and surface lightness is one property that aids recognition.

With this distinction in mind, let's take a look at some intriguing visual illusions that have been devised to test retinal mechanisms involving local spatial interactions.

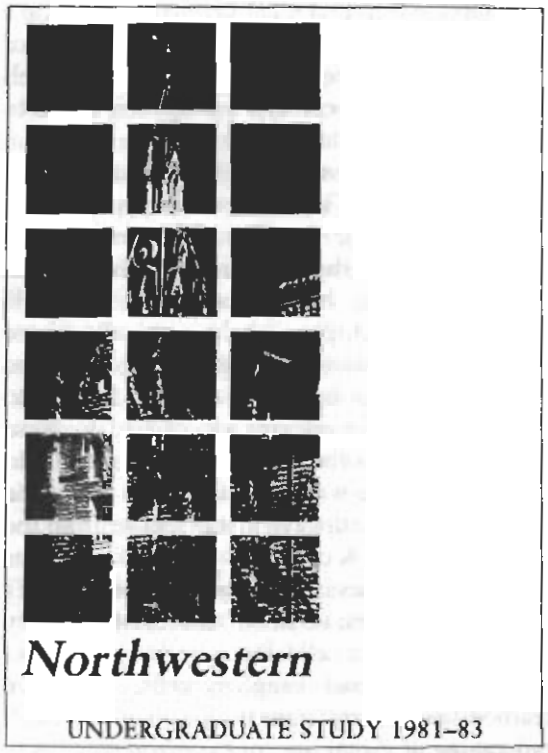

FIGURE 3.12
The cover of a catalog from one of America's outstanding institutions of higher education.

The Hermann Grid

Two decades ago, one of America's well-known private universities published what turned out to be a beguiling cover on the catalog of its course offerings (Figure 3.12). When the catalog was distributed, it generated a stir: viewers experienced what seemed to be spots that disappeared whenever they tried to look directly at them. By inspecting Figure 3.12, you can experience these illusory spots in most of the intersections between the horizontal and vertical white stripes. Note, too, that when you move your eyes to look *directly* at one of these phantom spots, it disappears.

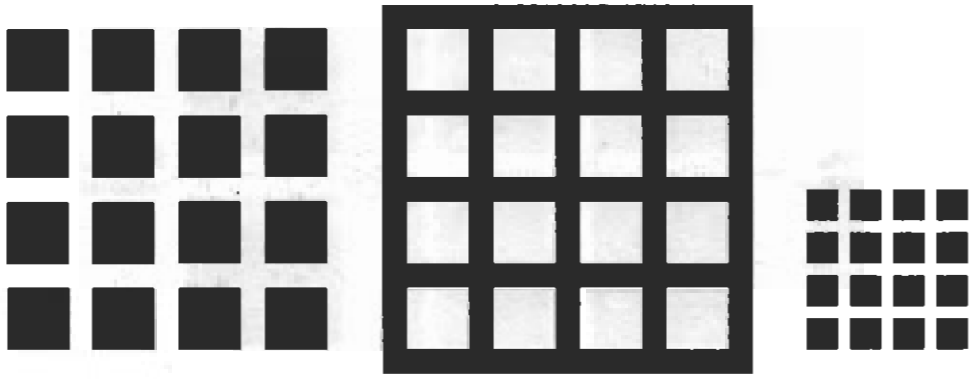

FIGURE 3.13
Three conventional versions of Hermann grids.

Actually, these illusory spots were first described more than a century ago. The pattern inadvertently used by the university resembled what is known as a **Hermann grid,** first described in the nineteenth century by Ludimar Hermann, a German physiologist. More traditional versions of Hermann's grid are shown in Figure 3.13. Looking at the left-hand grid, you'll notice dark spots located in most of the intersections of white horizontal and vertical stripes. Looking at the middle grid (a photographic negative of the other), you'll notice light spots in most of the intersections of black horizontal and vertical stripes. Looking at the right-hand grid (a smaller version of the one on the left), you'll see dark spots at *all* the intersections. Although they are vivid, every one of these spots is illusory. Where, then, do they come from? And why do they appear where they do? It is generally accepted that these illusory spots are the product of center/surround antagonism within receptive fields of retinal ganglion cells.

We'll make use of ON-center receptive fields to explain why spots are seen in the left-hand grid. Two questions must be answered. First, why are the spots that you see located only at intersections between horizontal and vertical stripes, not elsewhere? Second, why do you *not* see a spot located in an intersection when you look directly at it? To answer the first question, we've drawn receptive fields on Hermann's grid (Figure 3.14). This allows you to compare how the retinal image of the grid would affect the two receptive fields shown, one being stimulated by an intersection and the other stimulated by part of a white stripe that is not at an intersection.

To determine the response of either retinal ganglion cell, we must analyze how each of its components—the center and the surrounding area—would be affected by the grid pattern. Assume that a viewer gazes steadily on the spot labeled "fixation point." Now, note that the centers of both receptive fields receive the same amount of light, but the surrounds receive different amounts. Remember that light falling in an OFF portion of a receptive field reduces that cell's activity. This means that the cell whose receptive field is centered *on* the intersection will respond less than the cell whose receptive field is centered *between* intersections. Consequently, between intersections, the white stripes will look comparatively lighter. Since the reduced response is confined to cells with receptive fields centered on the intersections, one experiences dimming at such locales—gray spots.

Why, though, don't you see a spot in the intersection when you look directly at it? Recall that receptive fields vary in size according to their

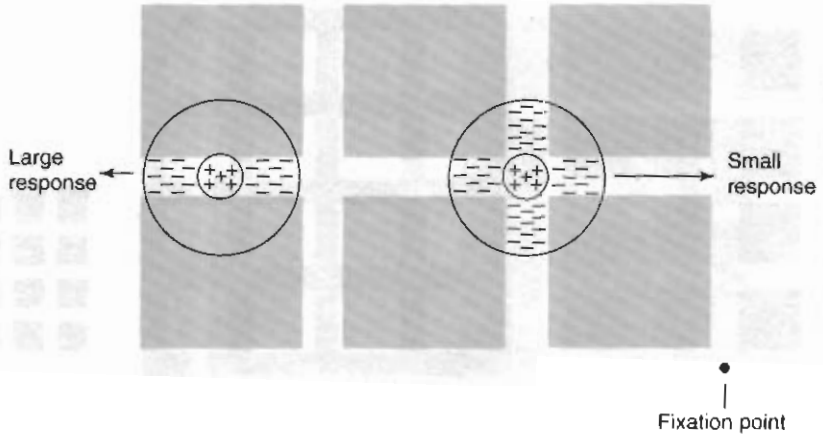

FIGURE 3.14
A possible neural explanation of Hermann grids.

eccentricity, the smallest receptive fields coinciding with the fovea. When you look directly at an intersection, you are using receptive fields whose centers and surrounds are so small that *both* fit completely within the width of a stripe. We've illustrated this in Figure 3.15.

Assume that a person fixates on the right-hand intersection of the grid. As you can see, these small receptive fields all receive the same amount of stimulation within their centers and surround. Consequently, all the cells around the region of fixation will give the same response, whether on the intersection or not. As a result, there will not be any local dimming at that intersection.

To test your understanding of these ideas, see if you can apply this same line of reasoning to the middle grid in Figure 3.13, where light illusory spots are seen. Finally, see if you can explain why spots are seen at *every* intersection, including the one you are fixating on, in the right-hand grid in Figure 3.13.

You might also try looking at Figure 3.13 under very dimly lit conditions. (Be patient; you must let your eyes adapt to the darkness.) You'll be amazed to see that the phantom spots, unlike most apparitions, actually disappear when the lights are low (Wist, 1976). Why should this happen? To give you a clue, remember that the illusory spots arise because of the surround portions of the receptive fields, which are differentially stimulated depending on what part of the grid they are analyzing. If we were to eliminate the surround's contribution to each cell's response, the spots should disappear. Evidently that's exactly what happens: the surround portions of ganglion cell receptive fields become ineffective at low-light levels, leaving only the center region to generate responses (Barlow, Fitzhugh and Kuffler, 1957).

Some features of the Hermann grid illusion aren't readily explained by the responses of ganglion cells (Wolfe, 1984). For one thing, the number of intersections present in the figure influences the strength of the illusory spots at those intersections. Increasing the number of intersections strengthens the illusion. The regularity of the intersection's spacing also modulates the strength of the illusion. These observations imply that far-reaching influences, not just local center/surround interactions, also play a role in creating these phantom spots.

Let's turn from illusory spots to another peculiarity of human vision that suggests the involvement of center/surround antagonism, Mach bands.

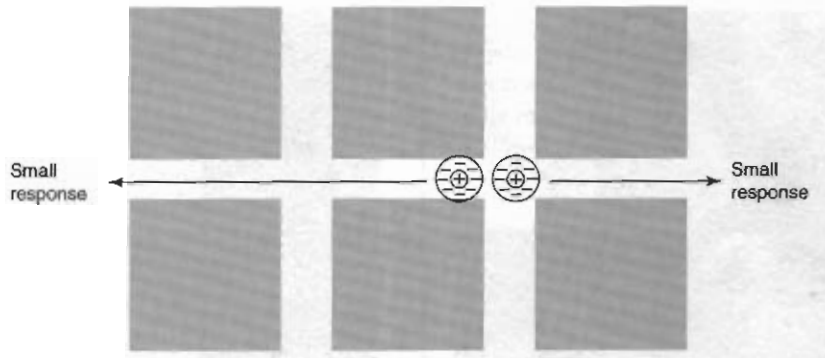

FIGURE 3.15
A possible neural explanation for the absence of local dimming of Hermann grids at the point of fixation.

Mach Bands

In Chapter 2 you were introduced to Ernst Mach, the Austrian physicist and philosopher who made important contributions to a number of scientific disciplines during the last part of the nineteenth century and the early part of the twentieth. We're concerned here with one small part of his work. The interested student will find an excellent, highly readable account of Mach's life and work in Ratliff (1965).

Mach became interested in the connection between the intensity of reflected light and the sensation it engendered. To explore this interest, Mach created various patterns out of paper that portrayed gradients of light, ranging from white through gray to black. Mach carefully studied whether perception of the gradients conformed to the actual distribution of light reflected from those patches of paper. In some instances, the two did not correspond. Mach had the great insight that these idiosyncrasies of perception were caused by antagonistic influences within the retina. Because we have much more information about retinal physiology and anatomy than was available to Mach, we are able to infer that these idiosyncrasies could be caused by the center/surround antagonism evidenced in the receptive fields of retinal ganglion cells.

The upper portion of Figure 3.16 shows one of the kinds of patterns Mach developed. The graph below the pattern plots the actual distribution of light intensity in the pattern. The horizontal axis of the graph represents position in the pattern. The vertical axis shows how much light the pattern reflects at that position. From left to right, the graph shows the level of intensity in the pattern increasing in a stepwise fashion. Thus, the pattern really consists of a number of bars, each of uniform intensity and each giving way abruptly to another level.

When you look at the pattern itself, though, you'll notice some things that don't seem to correspond with the graph. In particular, the lightness of each bar does not appear uniform. Take, for example, one of the bars in the middle. One edge—near the bar's left-hand, darker neighbor—seems extra light, whereas the other edge—near the bar's right-hand, lighter neighbor—seems extra dark. In other words, the *lightness* within the bar's interior varies even though the *intensity* of light reflected from the page does not. Most people describe the edges of each bar as having bands, extra dark and extra light regions. These bands are called **Mach bands** in honor of the first person to study them systematically.

These bands, no matter how vivid they may seem, are illusory. They don't exist on the paper,

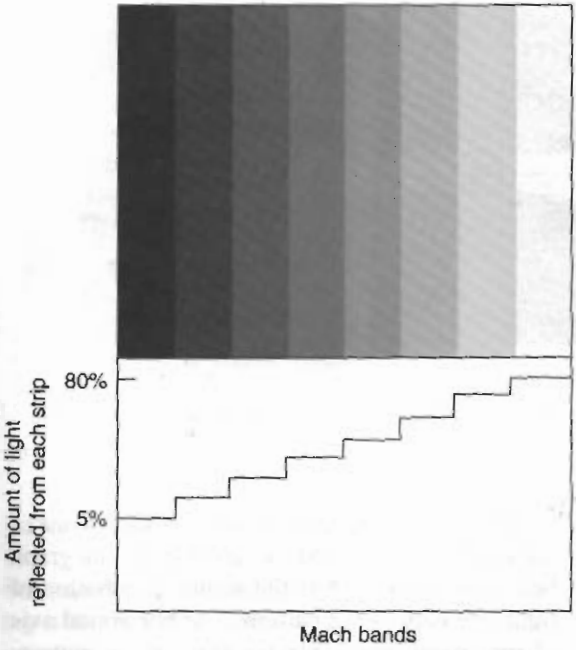

FIGURE 3.16
The lightness of each stripe in the pattern (upper portion of the figure) varies even though the intensity of each stripe is constant (lower portion of the figure).

only in your head. To understand the classic explanation of where they come from, consider how the retinal ganglion cells respond to such a pattern of bars. When one looks at the pattern in the upper portion of Figure 3.16, a distribution of light is produced on the retina similar to the distribution shown by the graph in the lower part of Figure 3.16. This distribution of light is broad enough so that it extends across the receptive fields of many retinal ganglion cells. To simplify the discussion, let's consider only three adjacent bars from the pattern (see Figure 3.17). Suppose that the image of the three bars falls on some small number, say five, of ON-center receptive fields (though this assumption is not crucial to our point). These ganglion cells are each responsible for signalling the brain about the intensity of light falling within their individual receptive fields. What sorts of messages would the brain

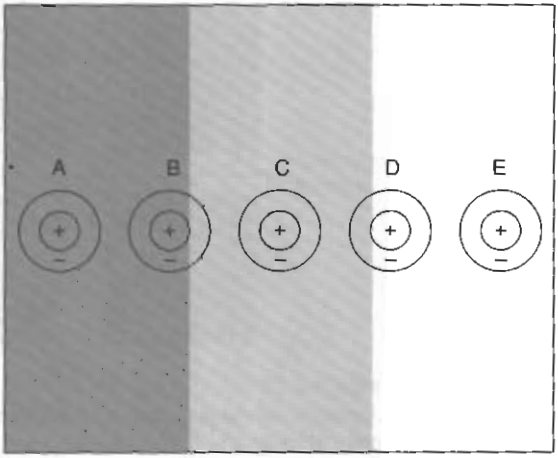

FIGURE 3.17
Possible neural explanation of Mach bands.

get if these cells were stimulated by the pattern shown in the upper part of Figure 3.16?

To see what response any cell would give, we must weigh how much light falls in each of its two regions. (Remember the convention: "plus" indicates an ON region, "minus" indicates an OFF region.) Let's take the three easiest ones first. Receptive field *A* (leftmost ganglion cell in Figure 3.17) receives the least light, field *E* the most, and field *C* an intermediate amount of light. Each cell's response will be proportional to the amount of light falling within its receptive field. So the bar *A* "sees" will seem the dimmest; the one *E* "sees," the brightest; and the one *C* "sees," appearing intermediate to the two. That leaves *B* and *D* as the interesting cases.

Receptive field *B*'s center is stimulated by the same level of light as *A*'s. Their respective surrounds, however, are *differentially* stimulated. All of *A*'s surround is dimly illuminated, thereby producing little antagonism to combat the response produced by the center. Although the left part of *B*'s surround is similarly illuminated, the right part is stimulated by the higher light level of the middle bar. As a result, the surround of *B* generates *more* antagonism than does the surround of

A, *diminishing* the overall response of *B* to a level below that of *A*. Consequently, the region *B* "sees" appears darker than *A* "sees." *B* creates a dark Mach band.

Now consider *D* and *E*. The net response from *D* will be larger than that from *E* because *D*'s surround is partially stimulated by the reduced light from the center bar, rather than by the higher level from the right-hand bar. As a result, *D*'s surround generates *less* antagonism to its center's response, yielding a net response that is *greater* than that from *E*. So the region *D* "sees" will appear lighter than that *E* "sees." *D* creates a light Mach band.

Mach bands emphasize the important distinction between intensity and lightness. **Intensity** is an objective, physical variable, something that a light meter can measure. **Lightness** is a subjective, perceptual variable, whose measurement requires a biological visual system. Often, intensity and lightness covary, with surfaces reflecting more intense light appearing lighter. But Mach bands show that the correlation is not perfect. Although intensity changes in a stepwise fashion, lightness does not.

Mach thought that these illusory bands arose from neural processing within a system that had center/surround antagonism. Numerous investigators since Mach have endorsed this idea and, going one step further, have identified that system as the retinal ganglion cells (for example, Cornsweet, 1970; Ratliff, 1984; Hubel, 1988; Dowling, 1998). Despite the plausibility of this simple, elegant theory, recent work has challenged the idea that center/surround antagonism in the retina provides an adequate account of lightness illusions such as Mach bands. Even when the pattern of retinal stimulation remains constant, the vividness of illusory Mach bands varies depending on how you interpret the surface upon which the intensity gradient appears (Lotto et al., 1999). If the gradient is construed to be a shadow on a flat surface, Mach bands are not nearly as salient as when that same gradient is associated with a curved surface (Figure 3.18). The same light intensity profile, in other words, is perceived differently simply by changing what it is you think you're looking at. There is no obvious reason why the responses of retinal ganglion cells in the eye should be susceptible to this kind of "cognitive" influence. Consequently, vision

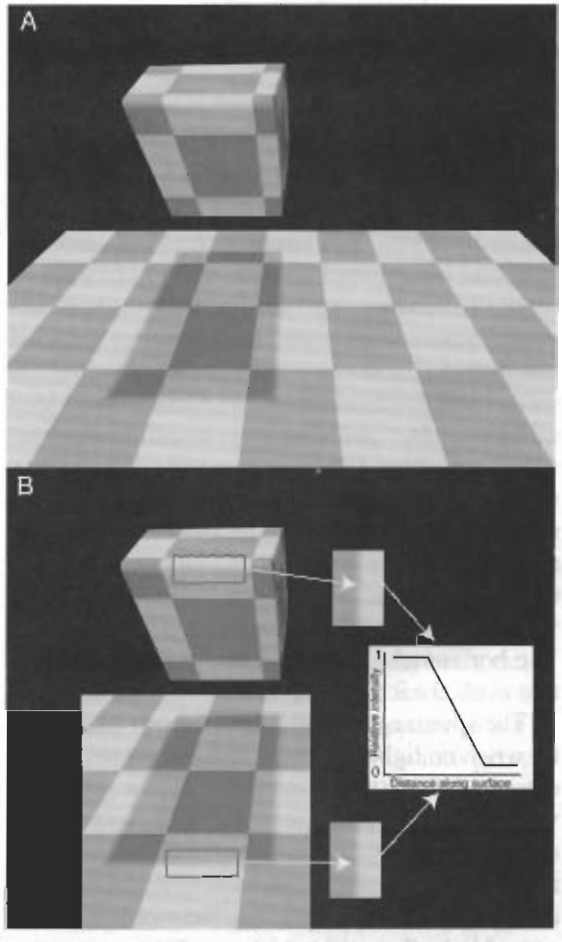

FIGURE 3.18
The vividness of Mach bands depends on the figure's interpretation. Looking at panel A, compare the lightness of the curved edge of the box (the upper outlined region in panel B) with the shadow cast on the textured floor (the lower outlined region in panel B). The two parts of the image are physically identical (as denoted by the graph in panel B), but the Mach band seems more pronounced in the region associated with the surface of the object than in the region of the shadow. (Lotto, Williams and Purves, 1999)

scientists now believe that while the center/surround antagonism in retinal ganglion cells does contribute to Mach bands, such antagonism cannot be the complete explanation.

Next let's consider another illusion of surface lightness, this one known as simultaneous lightness contrast. To set the stage for this illusion, let's perform a few more experiments on the retinal ganglion cells we've been talking about.

Two More Center/Surround Experminents

In the first experiment, we'll compare how a retinal ganglion cell responds when its center is stimulated by light of different intensities while its surround is subjected, in turn, to no illumination, some illumination, and finally, intense illumination. Suppose that we are recording from a retinal ganglion cell of the ON-center variety. Finding the cell's receptive field center, we focus a small spot of light within just that area, avoiding the surround altogether. The test spot is turned on for a second and the resulting number of impulses is recorded. After repeating this procedure with test spots of varying intensities, we plot the results as the curve labeled *A* in Figure 3.19. The graph's vertical axis represents the number of impulses. The horizontal axis represents the light intensity that evoked each response.

The spontaneous activity of the cell (the activity when no light is present at all) is plotted at the extreme left of the horizontal axis. The graph reveals that very weak intensities of light fail to change the cell's spontaneous activity. Notice the arrow extending from curve *A* to the horizontal axis of the figure. It indicates the weakest intensity of light that will evoke a response appreciably different from that occurring with no light at all. In a sense, this is the minimum amount of light that this cell can "distinguish" from complete darkness.

As the spot's intensity increases, the cell's response grows larger. Finally, at some high intensity, the response of the cell *saturates*, meaning that further increases in intensity fail to produce

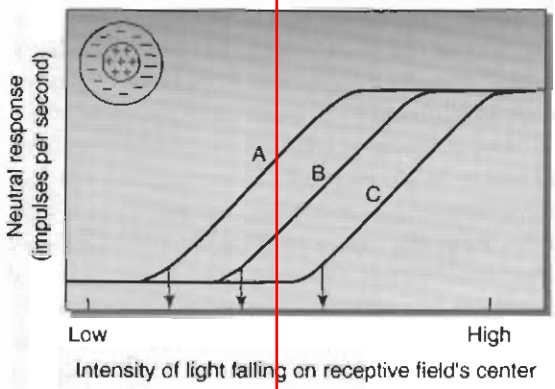

FIGURE 3.19
Graph showing the response of an ON-center ganglion cell to light of varying intensity falling on the center of the cell's receptive field. The three curves represent the effects of different intensities of surrounding stimulation.

corresponding increases in response. As far as this cell is concerned, all spots more intense than the point of saturation are indistinguishable from one another. From the cell's viewpoint, all such spots are equally intense.

Now suppose we repeat this experiment with one notable modification. The small spot again falls in the receptive field center, but now the surround is also illuminated. Again, we vary the intensity of the spot in the receptive field center, holding the surround illumination constant. The cell's response under this condition is shown by curve *B* in Figure 3.19. What are the differences between curves *A* and *B*? For both, arrows extending to the horizontal axis indicate the dimmest intensities that will produce an appreciable change in the cell's spontaneous response rate. As you can see, when the surround is illuminated, the center must be more strongly stimulated in order to change the cell's response. More generally, adding light in the surround reduces the response produced by any given intensity of light in the center. This is another demonstration of antagonistic forces at work: light has opposite effects on the ganglion cell's activity, depending on where the light strikes within the cell's receptive field.