

# CHAPTER 3



Jasper Johns, Flag, 1954–55

Art © Jasper Johns/Licensed by VAGA, New York, NY Photo credit: © The Museum of Modern Art/Licensed by SCALA/Art Resource, NY

© 2009 Sinauer Associates, Inc. This material cannot be copied, reproduced, manufactured, or disseminated in any form without express written permission from the publisher.

# Spatial Vision: From Stars to Stripes

© 2009 Sinauer Associates, Inc. This material cannot be copied, reproduced, manufactured, or disseminated in any form without express written permission from the publisher.

In Chapter 2 we learned that the macroscopic structures of the human eye function essentially as a biological camera: the iris regulates the number of light rays entering the eyeball; and the cornea, lens, and aqueous and vitreous humors focus these rays so that a clear image is formed on the retina. The rod and cone photoreceptors capture this image in a way that is roughly analogous to the way the film in a camera captures photographic images.

It is here, however, that the analogy between visual system and camera ends. Cameras take pictures. Visual systems see. How do we get from an image of the world in front of us to an interpretation of that world—what is out there, where it is, and what we can do to it? This process starts in the eyeball itself, where the postreceptor layers of the retina translate the raw light array captured by the photoreceptors into the patterns of spots surrounded by darkness, or vice versa, detected by the ganglion cells (see Figure 2.13). As we discussed in Chapter 2, this retinal translation helps us perceive the pattern of light and dark areas in the visual field regardless of the overall light level (e.g., it enables us to see almost as well at dusk as we can at noon).

In this chapter we follow the path of image processing from the eyeball to the brain (Figure 3.1). As we will see, neurons in the cerebral cortex translate the array of stars perceived by retinal ganglion cells into something like an array of stripes. Furthermore, we will see that this portion of visual cortex is organized into thousands of tiny computers, each responsible for determining the orientation, width, color, and other characteristics of the stripes in one small portion of the visual field. In Chapter 4, we will continue this story by examining how other parts of the brain assemble the outputs from these minicomputers to produce a coherent representation of the objects whose reflected light started the photoreceptors firing in the first place.

## Visual Acuity: Oh Say, Can You See?

*The King said, "I haven't sent the two Messengers, either. They're both gone to the town. Just look along the road, and tell me if you can see either of them."*

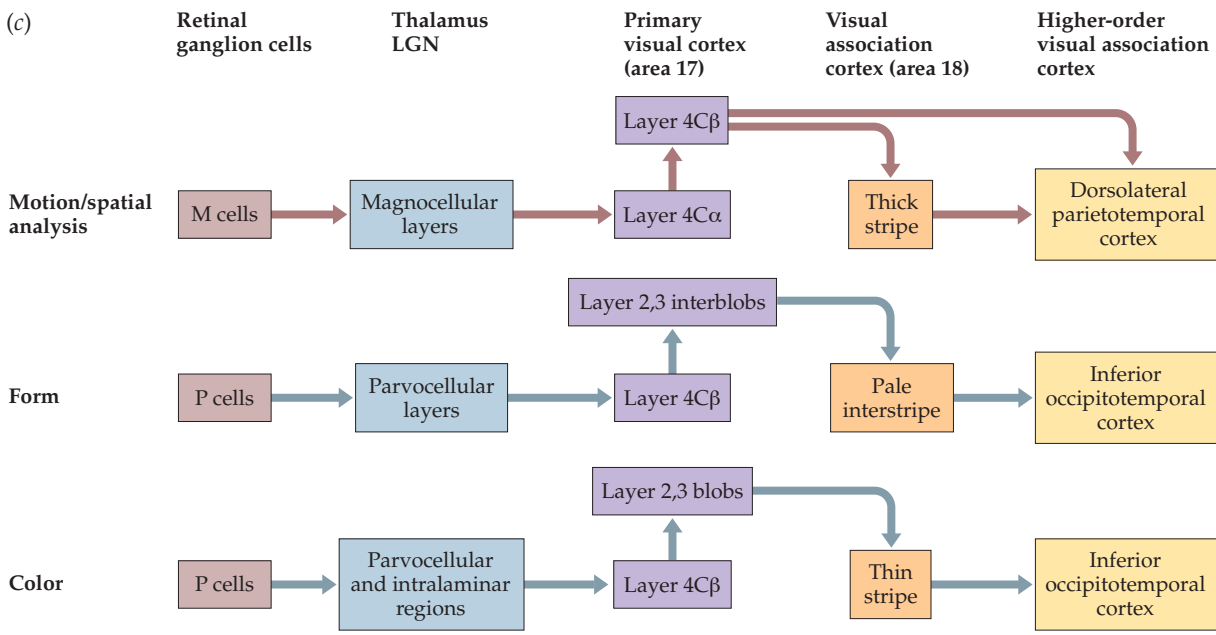
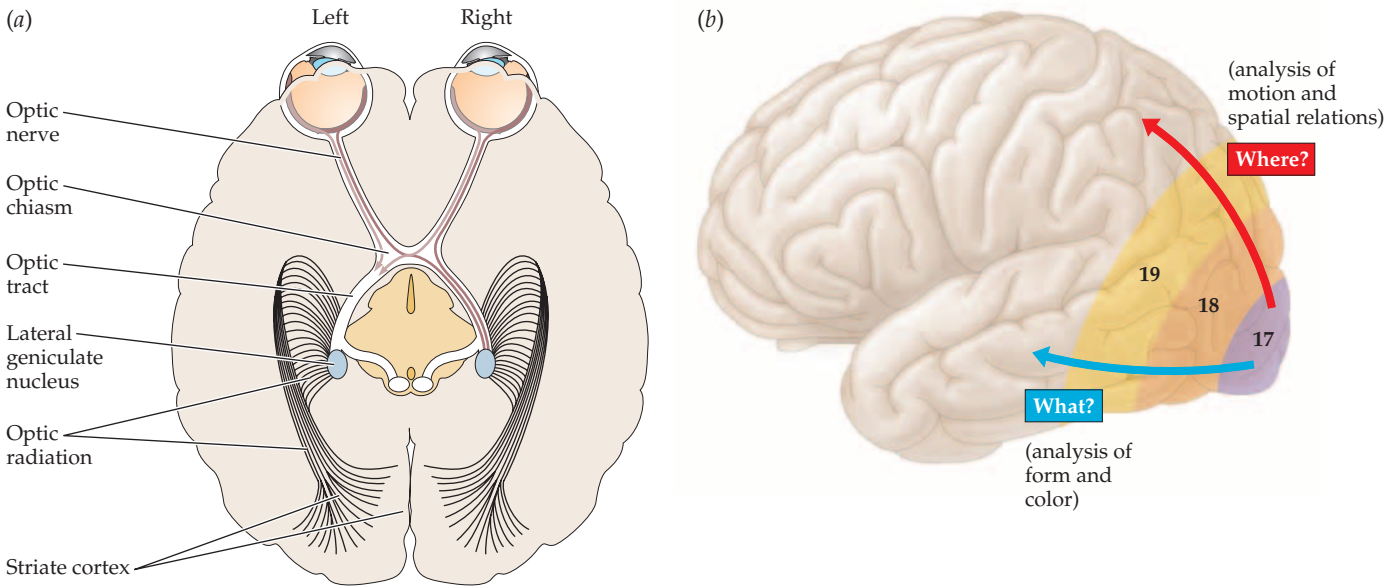
*"I see nobody on the road," said Alice.*

*"I only wish I had such eyes," the King remarked in a fretful tone. "To be able to see Nobody! And at that distance, too!"*

— Lewis Carroll, *Through the Looking Glass*

Since we'll be talking in this chapter about how the visual system codes images in terms of oriented stripes, let's start by determining just how well we see stripes when they are very close together and/or when the difference in illumination between the stripes and the background (the **contrast**) is very

**contrast** The difference in luminance between an object and the background, or between lighter and darker parts of the same object.




**FIGURE 3.1** Cortical visual pathways. (a) The basic organization of the primary visual pathway from eyeball to striate cortex, in coronal section. (b) A lateral section of the brain, illustrating early visual areas and showing the putative where (dorsal) and what (ventral) streams. (c) The flow of information for motion, form, and color analysis from retinal ganglion cells to higher-order visual association cortex. (Part a after Purves et al., 2008; b and c after Blumenfeld, 2002.)

**acuity** The smallest spatial detail that can be resolved.

low. In addition to setting the boundary conditions for how well we should expect the visual system to be able to perform, we will use this section to introduce some important jargon that we will need in the rest of the chapter.

Get a tape measure, prop your textbook up, and, while looking at the X in the middle of **Figure 3.2**, back up until you cannot tell the orientation of the black and white stripes. Measure how far your eye is from the page. Now walk forward a bit until you're sure you can see which grating includes vertical stripes and which horizontal stripes, and again measure your distance from the page. Congratulations! You just completed a fast (but not terribly accurate) measurement of your own visual resolution **acuity**. Eye doctors specify acuity in terms like 20/20 (more about this in a moment), but vision




**FIGURE 3.2** A visual acuity test. See the text for details.

scientists prefer to talk about the smallest visual angle of a cycle of the grating that we can perceive (**Figure 3.3**). A **cycle** is simply one repetition of a black and a white stripe (both of the gratings in **Figure 3.2** have 25 total cycles). **Visual angle**, which we alluded to in Chapter 2, is the angle that would be formed by lines going from the top and bottom (or left and right, depending on the orientation of the stripes) of a cycle on the page, through the center of the lens, and on to the retina. You can learn more about this concept in **Web Activity 3.1: Visual Angle**.


More precisely, to calculate the visual angle of your resolution acuity, divide the size of the cycle in **Figure 3.2** (which was 2 millimeters, or 1/16 inch) by the viewing distance at which you could just barely make out the orientation of the gratings (average your first and second measurements to get a rough estimate of this distance), and then take the arctangent of this ratio. Under ideal conditions, humans with very good vision can resolve gratings like those in **Figure 3.2** when one cycle subtends an angle of approximately 1 minute of arc (0.017 degree).

This resolution acuity represents one of the fundamental limits of spatial vision: it is the finest high-contrast detail that can be resolved. The limit is determined primarily by the spacing of photoreceptors in the retina. To see why, imagine that we're projecting the **sine wave gratings** shown in **Figure 3.4** onto the retina. The light intensity in such gratings varies smoothly and continuously across each cycle (unlike the gratings in **Figure 3.2**, in which intensity changes abruptly from black to white and back to black). However, the visual system "samples" the grating discretely, through the array of receptors at the back of the retina (in this respect the eye is more like a digital camera than like a traditional camera that uses film). If the receptors are spaced such that the whitest and blackest parts of the grating fall on separate cones (**Figure 3.4a**), we should be able to make out the grating. But if the entire cycle falls on a single cone (**Figure 3.4b**), we will see nothing but a gray field (or we may experience a phenomenon called **aliasing**, in which we misperceive the cycles to be longer than they actually are).



**FIGURE 3.3** Visual angle. Shown here is the angle size of one cycle of a grating at the retina.

**FIGURE 3.4** Sine wave gratings. (a) The stripes are wider than the photoreceptors (pink circles in the top panel), and the grating can be reconstructed vertically. (b) The stripes are narrower than the photoreceptors, so both black and white bars will fall inside a single receptor (top panel) resulting in a uniform gray field (bottom panel).



Cones in the fovea have a center-to-center separation of about 0.5 minute of arc (0.008 degree), which fits nicely with the observed acuity limit of 1 minute of arc (remember that we need two cones per cycle to be able to perceive it properly). Rods and cones in the periphery are packed together less tightly (recall that rods are physically more tightly packed [denser] than cones), but here many receptors converge on each ganglion cell. As a result, visual acuity is much poorer in the periphery than in the fovea. For a demonstration of the difference between foveal and peripheral vision, see [Web Activity 3.2: Foveal Acuity](#). (See also [Web Essay 3.1: Hyperacuity](#).)


### A Visit to the Eye Doctor

Eye doctors don't describe acuity in terms of visual angles and cycles. The last time you visited your eye doctor, she may have asked you to read letters, decreasing the size of the letters until you made several errors. Then she may have told you that your visual acuity was 20/20 if your vision was good, or 20/30 if you needed glasses, or possibly 20/10 if you could read the smallest letters on the eye chart. This strange method for designating visual acuity was invented in 1862 by a Dutch eye doctor, Herman Snellen (1834–1908). Snellen constructed a set of block letters for which the letter as a whole was five times as large as the strokes that formed the letter ([Figure 3.5](#)). Note that the resulting patterns are reminiscent of the gratings in [Figure 3.2](#). He then defined visual acuity as follows:

$$\frac{\text{(the distance at which a person can just identify the letters)}}{\text{(the distance at which a person with "normal" vision can just identify the letters)}}$$

In later adaptations of the Snellen test, the viewer was positioned at a constant distance of 20 feet, and the size of the letters, rather than the posi-

**FIGURE 3.5** A Snellen E. The letter size is five times the stroke size.



tion of the viewer, was altered. So normal vision came to be defined as 20/20. To relate this back to visual angle, a 20/20 letter is designed to subtend an angle of 5 arc minutes (0.083 degree) at the eye, and each stroke of a 20/20 letter subtends an angle of 1 arc minute (the familiar 0.017 degree). Thus, if you can read a 20/20 letter, you can discern detail that subtends 1 minute of arc. If you have to be at 20 feet to read a letter that someone with normal vision can read at 40 feet, you have 20/40 vision (worse than normal). Although 20/20 is often considered the gold standard, most healthy young adults have an acuity level closer to 20/15.

### Acuity for Low-Contrast Stripes

Up to now we've been discussing the tiniest high-contrast details that we can resolve. We learned that sine wave gratings in which the light stripes are perfectly white and the dark stripes perfectly black can be distinguished from a uniform gray field, as long as adjacent pairs of light or dark stripes are separated by at least 1 arc minute of visual angle. But what happens if the contrast of the stripes is reduced—that is, if the light stripes are made darker and the dark stripes lighter?


This was the question asked by Otto Schade in 1956, when he was working for the RCA Corporation. Schade showed people sine wave gratings with different spatial frequencies and had them adjust the contrast of the gratings until they could just be detected. **Spatial frequency** refers to the number of times a pattern, such as a sine wave grating, repeats in a given unit of space. For example, if you view your book from about 120 centimeters (about 47 inches) away, the visual angle between each pair of white stripes in **Figure 3.6b** is about 0.25 degree, so the spatial frequency of this grating is  $1/0.25 = 4$  **cycles per degree**. **Figure 3.6a** shows a grating with a relatively lower spatial frequency (about 2 cycles per degree), and **Figure 3.6c** illustrates a relatively higher spatial frequency (about 8 cycles per degree). **Web Activity 3.3: Gabor Patches** provides additional illustrations of sine wave gratings at different spatial frequencies.

Intuitively, you might think that the wider the stripes (that is, the lower the spatial frequency), the easier it would be to distinguish the light stripes from the dark stripes. But this is not exactly what Schade found. He, and later Fergus Campbell and Dan Green (1965), demonstrated that the human **contrast sensitivity function (CSF)** is shaped like an upside-down *U*, as shown in **Figure 3.7**. We obtain the units for the y-axis in this graph by taking the recip-

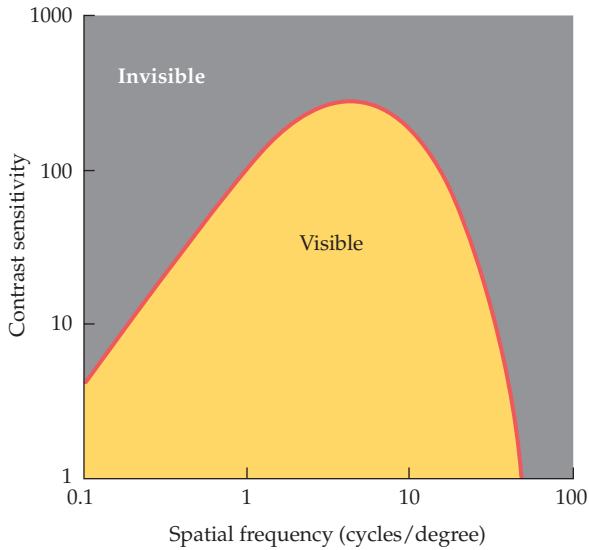
**spatial frequency** The number of cycles of a grating per unit of visual angle (usually specified in degrees).

**cycles per degree** The number of dark and bright bars per degree of visual angle.

**contrast sensitivity function (CSF)** A function describing how the sensitivity to contrast (defined as the reciprocal of the contrast threshold) depends on the spatial frequency (size) of the stimulus.



**FIGURE 3.6** Sine wave gratings illustrating low (a), medium (b), and high (c) spatial frequencies.




**FIGURE 3.7** The contrast sensitivity function (red line): our window of visibility. Any object whose spatial frequencies and contrasts fall within the yellow region will be visible. Those outside the yellow region are outside the window of visibility. The red line delimits the “threshold” between seeing and not seeing.

rocal of the **contrast threshold**. For example, for a 1-cycle/degree grating to be just distinguishable from uniform gray, the dark stripes must be about 1% darker than the light stripes (that is, if a tiny patch of a light stripe reflects 1000 photons, a patch of a dark stripe should reflect 990 photons). The reciprocal of this threshold is  $1/0.01 = 100$ , so this is the point plotted on the red CSF line in Figure 3.7 for this spatial frequency.

Note that a contrast of 100% corresponds to a sensitivity value of 1. The CSF reaches this value on the far right side of the curve in Figure 3.7, at about 60 cycles/degree. Sixty cycles/degree corresponds to a cycle width of 1 minute of arc, the resolution limit we measured previously for high-contrast stripes, which, recall, is determined primarily by cone spacing. The falloff in the CSF on the other side of the curve cannot be explained by cone spacing or by limitations in the optics of the eye. Instead, this part of the function must be due to neural factors, which we will discuss later.

You can visualize your own CSF by using **Figure 3.8**. Here we see a sinusoidal grating whose contrast increases continuously from the top of the figure to the bottom, and whose spatial frequency increases continuously from the left side of the graph to the right side. If you view the figure from a distance of about 2 meters, you will notice the inverted *U* shape where the grating fades from visibility to invisibility. If you bring the book closer to your eye, you should be able to see the stripes on the right side of the figure going farther up, whereas the tops of the stripes on the left side will become less distinct.

**contrast threshold** The smallest amount of contrast required to detect a pattern.



**FIGURE 3.8** A grating modulated by contrast (vertically) and by spatial frequency (horizontally). (From Robson and Campbell, 1997; courtesy of Izumi Ohzawa.)

## Why Sine Wave Gratings?

Now that we have described the contrast sensitivity function in some detail, one question may be foremost in your mind: Who cares? We don't see sine wave gratings in the real world; we see images of objects and scenes. What does contrast sensitivity for sine wave gratings across different spatial frequencies tell us about how we see real-world images?

One answer to this question is that, although "pure" sine wave gratings may be rare in the real world, patterns of stripes with more or less fuzzy boundaries are quite common: think of trees in a forest, books on a bookshelf, and a map of Manhattan (the latter includes a pattern of horizontal stripes superimposed on a pattern of vertical stripes). Furthermore, the edge of any object produces a single stripe, often blurred by a shadow, in the retinal image.


On a larger scale, the visual system appears to break down real-world images into a vast number of components, each of which is, essentially, a sine wave grating with a particular spatial frequency. This method of processing is analogous to the way in which the auditory system deals with sound and is called Fourier analysis (see Chapter 9). We'll return to this idea later in this chapter. For now, rest assured that scientists don't use sine wave gratings just because they are convenient to manipulate in experiments (although they do make very nice stimuli).

## Retinal Ganglion Cells and Stripes


In Chapter 2 we learned that retinal ganglion cells respond vigorously to spots of light. As it turns out, each ganglion cell also responds well to certain types of stripes or gratings. **Figure 3.9** shows how an ON retinal ganglion cell responds to gratings of different spatial frequencies. When the spatial frequency of the grating is too low (Figure 3.9a), the ganglion cell responds weakly because part of the fat, bright bar of the grating lands in the inhibitory surround, damping the cell's response. Similarly, when the spatial frequency is too high (Figure 3.9c), the ganglion cell responds weakly because both dark and bright stripes fall within the receptive-field center, washing out the response. But when the spatial frequency is just right (Figure 3.9b), with a bright bar filling the center and dark bars in the surround, the cell responds vigorously. Thus, retinal ganglion cells are "tuned" to spatial frequency: each cell responds best to a specific spatial frequency that matches its receptive-field size, and it responds less to both higher and lower spatial frequencies.

Christina Enroth-Cugell and John Robson (1984) were the first to record the responses of retinal ganglion cells to sinusoidal gratings. In addition to showing that these cells respond vigorously to gratings of just the right size, these investigators discovered that responses depend on the **phase** of the grating—its

**phase** The relative position of a grating.



**FIGURE 3.9** The response (right) of a ganglion cell to gratings of different spatial frequencies (left): (a) low, (b) medium, and (c) high.



**FIGURE 3.10** The response of a ganglion cell to a grating depends on the phase of the grating. This figure illustrates the response of an ON-center retinal ganglion cell to four different phases of an optimally sized grating.


position within the receptive field. **Figure 3.10** illustrates how an ON-center retinal ganglion cell might respond to a grating of just the right spatial frequency (a bar width about the size of the receptive-field center) in four different phases.

When the grating has a light bar filling the receptive-field center and the dark bars in the surround (Figure 3.10a), this ON-center cell responds vigorously, increasing its firing rate. If the grating phase is shifted by 90 degrees (Figure 3.10b), half the receptive-field center will be filled by a light bar and half by a dark bar, and similarly for the surround. In other words, there will be no net difference between the light intensity in the receptive field's center and its surround. In this case the cell's response rate does not change from its resting rate when the grating is turned on—just what we would predict if the ganglion cell were summing the total amount of light falling on its center and its surround. A second 90-degree shift puts the dark bar in the center and the light bars in the surround, producing a negative response (Figure 3.10c) (see Chapter 2); and a third phase shift returns us to the situation after the first shift, with the overall intensities in the center and surround equivalent and the cell therefore blind to the grating (Figure 3.10d). (Note that other ganglion cells respond to the 90-degree and 270-degree phases but not to the 0-degree and 180-degree phases, which is why the visual system as a whole is able to see all four phases equally well.)

**lateral geniculate nucleus (LGN)** A structure in the thalamus, part of the mid-brain, that receives input from the retinal ganglion cells and has input and output connections to the visual cortex.

**magnocellular layers** The neurons in the bottom two layers of the lateral geniculate nucleus, which are physically larger than those in the top four layers.

**parvocellular layers** The neurons in the top four layers of the LGN, which are physically smaller than those in the bottom two layers.



**FIGURE 3.11** The primate lateral geniculate nucleus. (From Hubel, 1988.)

### The Lateral Geniculate Nucleus

The axons of retinal ganglion cells synapse in the two **lateral geniculate nuclei (LGNs)**, which act as relay stations on the way from the retina to the cortex (see Figure 3.1). **Figure 3.11** shows that the LGN of primates is a six-layered structure, a bit like a stack of pancakes that has been bent in the middle (which is how the *geniculate*, which means “bent,” got its name). The neurons in the bottom two layers are physically larger than those in the top four layers; for this reason, the bottom two are called **magnocellular layers** and the top four **parvocellular layers** (*parvo-* is Latin for “small”). The two types of layers also differ in another, more important, way: the magnocellular layers receive input from M ganglion cells in the retina, and the parvocellular layers receive input from P ganglion cells. Functionally, studies in which magnocellular and parvocellular layers are chemi-

cally lesioned indicate that the magnocellular pathway responds to large, fast-moving objects, and the parvocellular pathway is responsible for processing details of stationary targets.

The organization of the retinal inputs to the LGN, diagrammed in **Figure 3.12**, provides some important insights into how our visual world is mapped to the brain. First, the left LGN receives projections from the left sides of the retinas in both eyes, and the right LGN receives projections from the right sides of the retinas. Second, each layer of the LGN receives input from one or the other eye. From bottom to top, layers 1, 4, and 6 of the right LGN listen to the left (**contralateral**) eye, while layers 2, 3, and 5 receive input from the right (**ipsilateral**) eye.


Each LGN layer contains a highly organized map of a complete half of the visual field. **Figure 3.12** shows schematically how objects in the right visual field (objects to the right of where our gaze is fixated) are mapped onto the different layers of the left LGN (as shown in **Figure 3.12**, the right side of the world falls on the left side of the retina, whose ganglion cells project to the left LGN). This ordered mapping of the world onto the visual nervous system, which is known as **topographical mapping**, provides us with a neural basis for knowing where things are in space (we will come back to this point a little later).

LGN neurons have concentric receptive fields that are very similar to those of retinal ganglion cells: they respond well to spots and gratings. Given that the LGN cells respond to the same patterns as the ganglion cells that provide their input, you might wonder why the visual system wastes a synapse here. That is, why don't the ganglion cell axons simply travel directly back to the cerebral cortex? One important reason probably has to do with the fact that there are more connections from other parts of the brain to the LGN than there are connections from the LGN to the cortex. In other words, the LGN appears to serve as a staging area where feedback from other parts of the brain modulates input from the eyes.

For example, the LGN is part of a larger brain structure called the "thalamus" (the medial geniculate nucleus, part of the auditory pathway, is another portion of the thalamus); and when we go to sleep, the entire thalamus is inhibited. Thus, even if someone pulls your eyelids open at night, you will not see anything (unless the physical contact awakens you). Input will travel from your retinas to your LGNs, but the neural signals will stop there before reaching the cortex, so they will never be registered. (The thalamic inhibition is not complete, however, which is why loud noises or very bright lights will be perceived, and will cause you to wake up and deal with the stimulus.)

### Striate Cortex

If you place one hand at the back of your head, about an inch or two above the top of your neck, you should be able to feel a small bump known as the "inion." The receiving area for LGN inputs in the cerebral cortex lies below the inion. This area has several names: **primary visual cortex, V1, area 17**, or **striate cortex** (because of the striped pattern it takes on after a certain type of staining procedure). By now you are probably getting the idea that layers are an important property of neural structures in the visual pathway. The striate cortex consists of six major layers, some of which have sublayers (**Figure 3.13**). Fibers from the LGN project mainly (but not exclusively) to layer 4, with magnocellular axons coming in to sublayer 4C $\alpha$  and parvocellular axons projecting to sublayer 4C $\beta$ . (See **Web Essay 3.2: The Whole Brain Atlas**.)



**FIGURE 3.12** Input (in this case the letters ABCDEF) from the right visual field is mapped in an orderly fashion onto the different layers of the left LGN, and input from the left visual field is mapped to the right LGN. Information from the two eyes is segregated into separate layers. Layers 1 and 2 are the magnocellular layers; layers 3 through 6 are the parvocellular layers.


**contralateral** Referring to the opposite side of the body (or brain).

**ipsilateral** Referring to the same side of the body (or brain).

**topographical mapping** The orderly mapping of the world in the lateral geniculate nucleus and the visual cortex.

**primary visual cortex (V1)** The area of the cerebral cortex of the brain that receives direct inputs from the lateral geniculate nucleus, as well as feedback from other brain areas, and is responsible for processing visual information. Also called area 17 or striate cortex.

**FIGURE 3.13** Striate cortex. (From Hubel, 1988.)



Like the LGN, the striate cortex has a systematic topographical mapping of the visual field. But the striate cortex is not simply a bigger version of the LGN. A major and complex transformation of visual information takes place in the striate cortex. For starters, striate cortex contains on the order of 200 million cells—more than 100 times as many as the LGN has!

### Cortical Topography and Cortical Magnification


**Figure 3.14** illustrates two important features of the visual cortex. The first is the topographical mapping already noted. The fact that the image of the woman's right eyebrow (*her* right; it appears on the left in Figure 3.14) is mapped onto regions corresponding to the numbers 3 and 4 in the striate cortex, tells the visual system that the eyebrow must be in positions 3 and 4 of the visual field.

The second feature evident in Figure 3.14 is the dramatic scaling of information from different parts of the visual field. Objects imaged on or near the fovea are processed by neurons in a large part of the striate cortex, but objects imaged in the far right or left periphery are allocated only a tiny portion of the striate cortex. This distortion of the visual-field map on the cortex is known as **cortical magnification** because the cortical representation of the fovea is greatly magnified compared to the cortical representation of peripheral vision. (See **Web Essay 3.3: Seeing Images on the Cortex.**)

To gain a sense of the extent of this cortical magnification factor, hold your two arms out in front of you, put up your index fingers, hold them about 10 centimeters (4 inches) apart, close your left eye, and fixate your right finger. In this position, your right fingernail, which is taking up about 1 degree of visual angle on the fovea, is being processed by neurons in about 20 millimeters (mm) of striate cortex. Your left fingernail, which is covering the same amount of visual angle but is falling 10 degrees to the left of the fovea, is being processed by only 1.5 mm of cortex.

One important consequence of cortical magnification is that visual acuity declines in an orderly fashion with eccentricity (distance from the fovea) (Levi, Klein, and Aitsebaomo, 1985)—a phenomenon demonstrated by Hermann Rudolf Aubert well over a century ago (Aubert, 1886). **Web Activity 3.2: Foveal Acuity** allows you to demonstrate this phenomenon yourself, as

**cortical magnification** The amount of cortical area (usually specified in millimeters) devoted to a specific region (e.g., 1 degree) in the visual field.



**FIGURE 3.14** The mapping of objects in space onto the visual cortex. This figure illustrates both the topographical mapping and the dramatic magnification of the foveal representation in the cortex. (After Frisby, 1980.)

© 2009 Sinauer Associates, Inc. This material cannot be copied, reproduced, manufactured, or disseminated in any form without express written permission from the publisher.

does **Figure 3.15**, in which the letters are scaled in size such that each one covers an approximately equal cortical area.

Why is the foveal representation in the cortex so highly magnified? The visual system must make a trade-off. High resolution requires a great number of resources: a dense array of photoreceptors, one-to-one lines from photoreceptors to retinal ganglion cells, and a large chunk of striate cortex (not to mention the real estate in other areas of cortex necessary to do something with the visual information coming out of V1). To see the entire visual field with such high resolution, we might need eyes and brains too large to fit in



**FIGURE 3.15** A letter chart in which the letter size increases with eccentricity in proportion to the inverse cortical magnification factor. If you fixate your gaze on the far left side of the figure, all seven letters should be equally easy to see.

our heads! Thus, we have evolved a visual system that provides high resolution in the center and lower resolution in the periphery. If you need to process the details of an object in the corner of your eye, you can simply turn your eye or head so that the object falls on the fovea instead.

### Receptive Fields in Striate Cortex

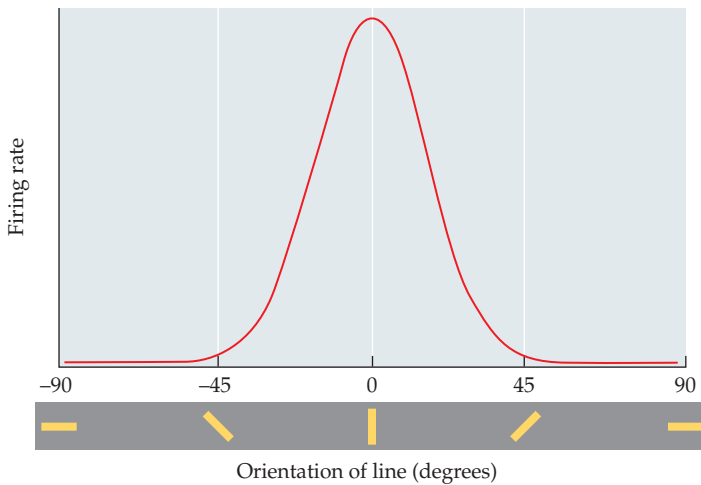
In 1958, David Hubel and Torsten Wiesel began work as postdoctoral students in Stephen Kuffler's laboratory. Their goal was to extend Kuffler's groundbreaking work on retinal ganglion cells to the cortex. So they began trying to map the receptive fields of neurons in striate cortex using spots of light, much as Kuffler (1953) had done earlier (see Chapter 2). To their dismay, however, they found that the cat's cortical cells hardly responded at all to the same spots that made its ganglion cells fire like crazy. To project their stimuli onto the retina, Hubel and Wiesel inserted a glass slide with a black spot into a slot in a special ophthalmoscope (that's the bright light the eye doctor shines in your eye in order to see your retina). One day, they had been recording from a neuron without much luck, when suddenly the cell emitted a strong burst of firing as they inserted the glass slide into the slot. Eventually they realized that the response had nothing to do with the spot itself; instead, the cell had been responding to the shadow cast by the edge of the glass slide as it swept across the ophthalmoscope's light path. And the rest, as they say, is history. Hubel related this story when he and Wiesel received the 1981 Nobel Prize in Physiology or Medicine for uncovering many of the remarkable properties of the visual cortex.

Hubel and Wiesel's most fundamental discovery was that neurons in striate cortex respond to stripes, not stars. More precisely, the receptive fields of striate cortex neurons are not circular, as they are in the retina and LGN. Rather, they are elongated, responding much more vigorously to bars, lines, edges, and gratings than to round circles of light.

### Orientation Selectivity

Further investigation by Hubel and Wiesel (1962) uncovered a number of other important properties of the receptive fields of neurons in striate cortex. First, an individual neuron will not respond equivalently to just any old stripe in its receptive field. It responds best when the line or edge is at just the right orientation, and hardly at all when the line is tilted more than 30 degrees away from the optimal orientation (a change equivalent to movement of the minute hand of a clock from 12 to 1). Scientists call this selective responsiveness **orientation tuning**: the cell is tuned to detect lines in a specific orientation in the same way that a piano key is tuned to produce a specific musical note.

**orientation tuning** The tendency of neurons in striate cortex to respond optimally to certain orientations, and less to others.



**FIGURE 3.16** Orientation tuning function of a cortical cell. The neuron fires vigorously when the line is oriented vertically, but it fires hardly at all when the line orientation is changed by 30 degrees.

A typical orientation tuning function looks like the plot in **Figure 3.16**. The neuron featured here fires vigorously when the line is oriented vertically (at 90 degrees), but the response tapers off rapidly as the line is tilted one way or another, diminishing to close to the cell's resting rate when the line is tilted 45 degrees in either direction. Other cells in striate cortex are selective for 0 degrees (horizontal), 45 degrees, 20 degrees, 62 degrees, and so on, so that the population of neurons as a whole detects all possible orientations. However, more cells are responsive to horizontal and vertical orientations than to obliques (De Valois, Yund, and Hepler, 1982; Li, Peterson, and Freeman, 2003). This physiological finding meshes well with the psychophysical finding that humans have somewhat lower visual acuity and contrast sensitivity for oblique targets than for horizontal and vertical targets.

How are the circular receptive fields in the LGN transformed into the elongated receptive fields in striate cortex? Hubel and Wiesel suggested a very simple scheme to accomplish this transformation (**Figure 3.17**). Simply put, their idea was that the concentric LGN cells that feed into a cortical cell are all in a row. Later studies (e.g., J. S. Anderson et al., 2000) have shown that the arrangement of LGN inputs is indeed crucial for establishing the orientation selectivity of striate cortex cells. However, other evidence suggests that neural interactions (e.g., lateral inhibition) within the cortex also play an important role in the dynamics of orientation tuning (Pugh et al., 2000).


### Other Receptive-Field Properties

Cortical cells respond not just to bars, lines, and edges. Like retinal ganglion cells, they also respond well to gratings (which are, after all, collections of lines). And like ganglion cells, they respond best to gratings that have just the right spatial frequency to fill the receptive-field center. That is, each striate cortex cell is tuned to a particular spatial frequency, which corresponds to a particular line width. Indeed, cortical cells are much more narrowly tuned (they respond to a smaller range of spatial frequencies) than retinal ganglion cells (De Valois, Albrecht, and Thorell, 1982). These narrow tuning functions mean that each striate cortex neuron functions as a **filter** for the portion of the image that excites the cell. We will return to the idea of striate cortex as a collection of filters later.


Another important discovery made by Hubel and Wiesel was that many cortical cells respond especially well to *moving* lines, bars, edges, and gratings. Moreover, many neurons respond strongly when a line moves in one

**filter** An acoustic, electrical, electronic, or optical device, instrument, computer program, or neuron that allows the passage of some frequencies or digital elements and blocks the passage of others.

**FIGURE 3.17** Hubel and Wiesel's model of how cortical simple cells get their orientation tuning.



**FIGURE 3.18** Drawing by René Descartes, 1664, illustrating the “fusion center.” (From Descartes, 1664.)



**ocular dominance** The property of the receptive fields of striate cortex neurons by which they demonstrate a preference, responding somewhat more rapidly when a stimulus is presented in one eye than when it is presented in the other.

**simple cell** A cortical neuron with clearly defined excitatory and inhibitory regions.


**complex cell** A neuron whose receptive-field characteristics cannot be easily predicted by mapping with spots of light.

direction—say, from left to right—but not at all when the same line moves from right to left.


As noted earlier, information from the two eyes is kept separate in the LGN: each LGN cell responds to one eye or the other, but never to both eyes. This arrangement changes dramatically in striate cortex, where a majority of cells can be influenced by input from both the left eye and the right eye. In other words, if a striate cortex neuron responds best to a 5-cycle/degree grating oriented at 45 degrees, it will respond to such a stimulus whether that stimulus is presented in the right eye or the left eye. However, striate cortex neurons often have a preference, responding somewhat more rapidly when a stimulus is presented in one eye than when it is presented in the other. Hubel and Wiesel called this property of striate receptive fields **ocular dominance**.

Given that we see a single, unified world, intuitively it makes sense that information from the two eyes should be brought together at some point. Until Hubel and Wiesel’s discovery, however, there were heated arguments about whether the information converged at all, and if so, whether it was in a specialized “fusion center” in the brain—a notion that dates back to Descartes (1664; see Howard and Rogers, 2001) (Figure 3.18). We’ll describe some of these debates in Chapter 6, when we discuss binocular vision.

(a) Edge detector



(b) Stripe detector




### Simple and Complex Cells

Like precortical neurons, cortical neurons come in a wide variety of types. Hubel and Wiesel characterized neurons as **simple cells** when those neurons had clearly defined excitatory and inhibitory regions. Figure 3.19 shows two varieties of simple-cell receptive fields and their preferred stimuli. An edge detector (see Figure 3.19a) prefers to see light on one side of its receptive field and darkness on the other side. A stripe detector (see Figure 3.19b) responds best to a line of light that has a particular width, surrounded on both sides by darkness. If a grating with the appropriate spatial frequency drifts across the receptive field of this cell, the cell’s response will be modulated as dark and bright bars drift across the receptive-field center, in exactly the same way that the response of the retinal ganglion cell shown in Figure 3.10 is modulated.

Other neurons show responses that cannot be simply predicted from their responses to stationary bars of light. Hubel and Wiesel called these **complex cells**. Like simple cells, each complex cell is tuned to a particular orientation

**FIGURE 3.19** Two flavors of simple cells: (a) an edge detector and (b) a stripe detector.




**FIGURE 3.20** A simple cell and a complex cell might both be tuned to the same orientation and stripe width (spatial frequency), but the complex cell will respond to that stripe presented anywhere within its receptive field, whereas the simple cell might respond to the stripe in only one position.

and spatial frequency and shows an ocular preference. However, whereas a simple cell might respond only if a stripe is presented in the center of its receptive field, a complex cell will respond regardless of where the stripe is presented, as long as it is somewhere within the cell's receptive field (**Figure 3.20**). When tested with a drifting grating, the complex cell gives a robust response, with little or none of the modulation shown by simple cells (as well as retinal ganglion and LGN cells). Another way of stating this difference is to say that simple cells are "phase-sensitive," and complex cells are "phase-insensitive."

As with all other neurons in the visual system, save retinal photoreceptors, the receptive fields of complex cells represent a pooling of the responses of several subunits. The subunits give the complex cell its spatial frequency and orientation tuning, but the complex pooling operation makes the complex cell insensitive to the precise position of the stimulus within its receptive field. Hubel and Wiesel hypothesized a hierarchy in which LGN cells fed into simple cells, which in turn provided excitatory inputs to complex cells. However, substantial evidence now suggests that complex cells represent a separate parallel pathway (that is, that both simple and complex cells get direct input from LGN neurons).

### Further Complications

Hubel and Wiesel described another property of some cells in striate cortex that they called **end stopping**. When they tested an end-stopped cell with bars of increasing lengths, the response rate first increased as the bar filled up its receptive field, and then decreased markedly as the bar was lengthened further (**Figure 3.21**). Hubel and Wiesel called these cells "hypercomplex" cells, although they now appear to be subclasses of the simple and complex cells already discussed here (that is, there are simple end-stopped cells and complex end-stopped cells). End stopping is thought to play an important role in our ability to detect luminance boundaries and discontinuities.



**FIGURE 3.21** When the stimulus (bar) does not reach the outside edge of the receptive field or extends beyond the receptive field of an end-stopped cortical neuron, the neuron fires less than when the stimulus is just the right length.

Recent research has revealed further idiosyncrasies in the receptive fields of striate cortex neurons. For example, the size of a particular cell's receptive field appears to vary with target contrast; for instance, the cell might respond to a smaller portion of the visual field when the grating stimulus has a high contrast as opposed to when the difference between light and dark bars is more subtle (Sceniak et al., 1999). And neurons can be influenced by stimuli that fall outside the classic receptive field, via short- or long-range lateral connections and/or via feedback from neurons in other layers (Zipsper, Lamme, and Schiller, 1996).

As is the case for most of the visual system, what we don't know about the workings of striate cortex neurons almost certainly dwarfs what we do know. But to review what we do know, spend some time with **Web Activity 3.4: Striate Receptive Fields**, where you can try your hand at determining the receptive field of an unknown virtual neuron.

## Columns and Hypercolumns

As we've discussed, each of the approximately 200 million neurons in striate cortex responds to a distinctive set of stimulus properties: stripes, edges, and/or gratings that are oriented at a particular angle, with a particular width or spatial frequency, possibly moving in a particular direction. Some neurons are simple cells and some are complex cells, and each one is end-stopped or not. Most neurons also respond preferentially to stimuli presented in one eye or another. And each neuron responds only when its preferred stimulus is presented in one particular part of the visual field.


Hubel and Wiesel noticed very early on that these various receptive-field properties are not scattered haphazardly around striate cortex. Once they had figured out what the cells were looking for (stripes, rather than spots), they discovered that if they pushed the recording electrode down through the layers of the cortex in a direction perpendicular to the cortical surface, all the cells they encountered showed similar orientation preferences. If they shifted the electrode position over a tiny distance and made another perpendicular penetration, all the cells now responded best to a slightly different orientation, perhaps 10 or 15 degrees from the original orientation. On the basis of these observations, Hubel and Wiesel concluded that neurons with similar orientation preferences are arranged in **columns** that extend vertically through the cortex.

When Hubel and Wiesel made tangential penetrations into striate cortex (inserting an electrode such that it was oriented parallel to the cortical surface, rather than perpendicular), they found a systematic and progressive change in preferred orientation so that essentially all the orientations were encountered in a distance of about 0.5 mm. This finding has been confirmed via alternative physiological techniques. **Figure 3.22a** shows a small portion of a monkey's striate cortex prepared so that neurons responding to vertically oriented lines are stained black, while other neurons remain white. The distance between the vertical orientation columns revealed by this technique is, sure enough, just about 0.5 mm (LeVay, Hubel, and Wiesel, 1975). The mapping of orientation in the cortex can be appreciated from optical imaging studies (**Figure 3.22c**), where orientation preference is indicated by color.


Orientation is not the only property arranged in columns in the visual cortex. Neurons that share the same eye preference (exhibiting what is referred to as "ocular dominance") also have a columnar arrangement (**Figure 3.22b**). Furthermore, single-cell recording experiments, as well as staining experiments, indicate that eye preference switches (you guessed it) every 0.5 mm or so.

**column** A vertical arrangement of neurons.


(a) Orientation columns



(b) Ocular dominance columns



(c)




**FIGURE 3.22** Orientation (a) and ocular dominance (b) columns of the striate cortex, revealed by staining. (c) Optical imaging of the orientation maps in monkey cortex. (a and b from Hubel et al., 1978; c from Nauhaus et al., 2008.)

Through their studies, Hubel and Wiesel arrived at the model of striate cortical architecture illustrated in **Figure 3.23**. They proposed that a 1-mm block of striate cortex contains “all the machinery necessary to look after everything the visual cortex is responsible for, in a certain small part of the visual world” (Hubel, 1982). Each of these sections of cortex is called a **hypercolumn**. It contains at least two sets of columns, each covering every possible orientation (0–180 degrees), with one set preferring input from the left eye and one set preferring input from the right eye.

Hypercolumns are roughly 1 mm across throughout the striate cortex, but because of the cortical magnification factor discussed earlier, not all hypercolumns see the world at the same level of detail. A hypercolumn in the part of the cortex that represents the fovea may “see” a portion of the visual field that is 0.05 degree of visual angle across; a hypercolumn responding to input 10 degrees to the right of the fovea should cover about 14 times as large an area (0.7 degree across).

**hypercolumn** A 1-millimeter block of striate cortex containing two sets of columns, each covering every possible orientation (0–180 degrees), with one set preferring input from the left eye and one set preferring input from the right eye.


**FIGURE 3.23** This model of a hypercolumn shows two ocular dominance columns (one for each eye) and many orientation columns, and illustrates the locations of the cytochrome oxidase (CO) blobs. (From Breedlove et al., 2007.)



**cytochrome oxidase (CO)** An enzyme used to reveal the regular array of “CO blobs,” which are spaced about 0.5 millimeter apart in the primary visual cortex.

Orientation and ocular dominance are probably not the only stimulus dimensions that have a systematic columnar arrangement in the visual cortex. For example, another staining technique, which takes advantage of an enzyme called **cytochrome oxidase (CO)**, has revealed a regular array of “CO blobs” (shown in section in **Figure 3.24**), spaced that magical distance of about 0.5 mm apart (see **Figure 3.23**). The functional role of these blobs remains unclear, but CO blob columns have been implicated in processing color, with the interblob regions (note the elegant scientific jargon that has developed around this field of study) processing motion and spatial structure (Livingstone and Hubel, 1988). This view is probably too simplistic, but the blob array does suggest some kind of additional organizational layer on top of the orientation and ocular dominance arrays.

In sum, the current state of understanding is that striate cortex is concerned with analyzing the orientation, size, shape, speed, and direction of motion of objects in the world, and that it does so using modular groups of neurons—hypercolumns—each of which receives input from and processes a small piece of the visual world. (You can explore these organizational principles interactively in **Web Activity 3.5: Hypercolumns**.) Combining information from multiple hypercolumns is presumably the job of other portions of cortex farther downstream in the visual system. We will consider some of these portions in Chapter 4, when we discuss the representation and recognition of whole objects.



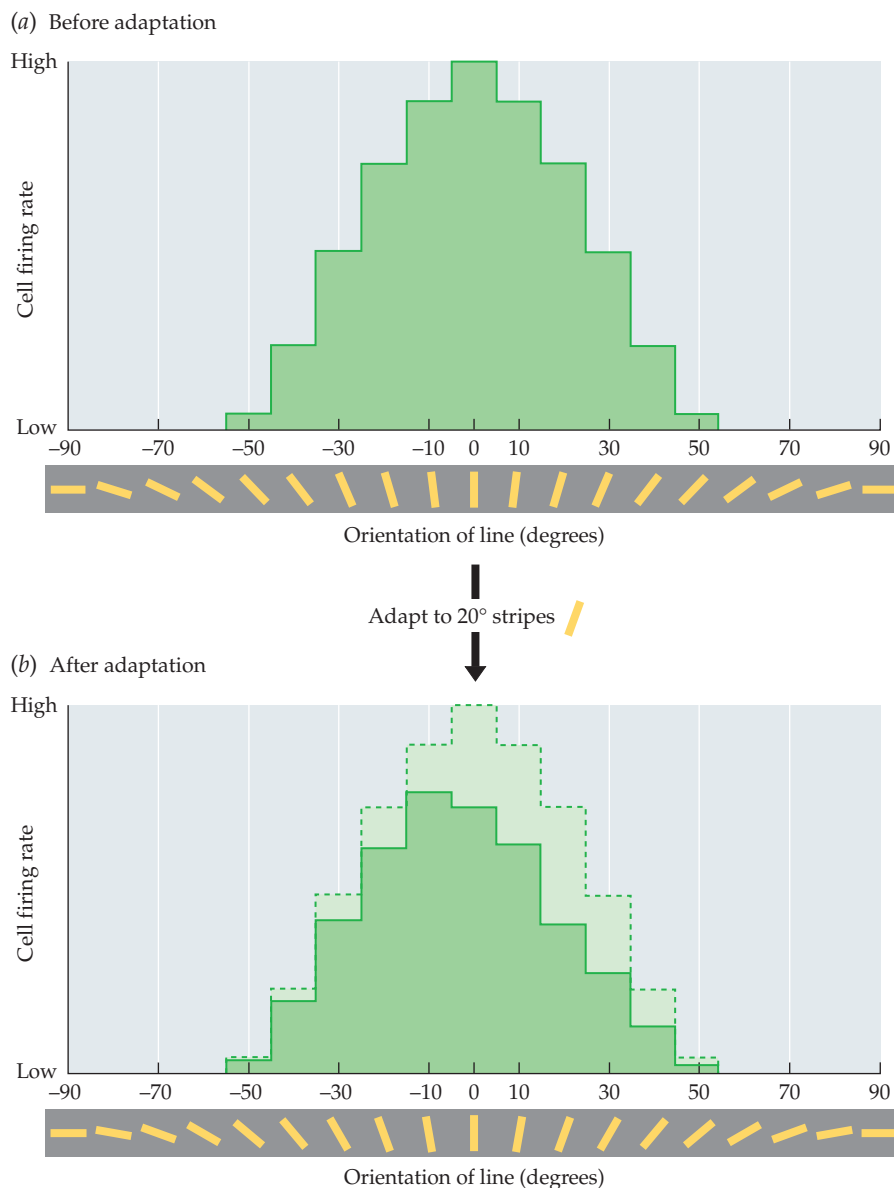
**FIGURE 3.24** Cytochrome oxidase (CO) blobs. (From Hubel, 1988.)

## Selective Adaptation: The Psychologist's Electrode

Most of the physiological research reported up to now in this chapter was done using cats, monkeys, or other animals as subjects. Does the human visual system also include neurons selective for orientation, line width, direction of motion, and so on? The difficult thing about answering this question is that we can't normally poke electrodes into a human's brain (which is why Hubel, Wiesel, and their peers had to use cats and monkeys in the first place), so indirect methods of learning about brain function had to be devised. One such method is called **adaptation**, a technique that gives psychologists a non-invasive "electrode" they can use to probe the human brain (Frisby, 1980).

Selective adaptation can provide insights into the properties of cortical neurons, as illustrated in **Figure 3.25**. The green bars in **Figure 3.25a** illustrate the normal firing rates of cells tuned to orientations of 0, 10, -10, 20, -20 degrees, and so on, to a vertical grating. By definition, gratings oriented at 0

**adaptation** A reduction in response caused by prior or continuing stimulation.




**FIGURE 3.25** The psychologist's electrode. This schematic diagram shows how selective adaptation may alter the distribution of neural responses and therefore perception. See the text for explanation.

**tilt aftereffect** The perceptual illusion of tilt, produced by adaptation to a pattern of a given orientation.

degrees (vertical) elicit the strongest response from the 0-degree cells, followed closely by the 10-degree and  $-10$ -degree cells, followed by the 20-degree and  $-20$ -degree cells, and so on. Now suppose we expose the visual system that contains these cells to a 20-degree grating for an extended period of time. This adapting stimulus will cause the 20-degree-selective cells to be most active, and the extended activity will fatigue these cells (that is, their maximum firing rate will be reduced for a short period following adaptation). The adaptation procedure will also affect the other cells to some extent: the 10-degree and 30-degree cells will be the next most fatigued, followed by the 0-degree and 40-degree cells, and so on.

Figure 3.25*b* shows what should happen when we present the vertical grating again after adaptation to the 20-degree grating, assuming that our orientation perception is really due to populations of orientation-selective cells like those that Hubel and Wiesel found in the cat cortex. As the dark green bars show, because the 0-degree cells have been fatigued more than the  $-10$ -degree cells, the  $-10$ -degree cells are now firing fastest. As a result, we should perceive the vertical test stimulus as being oriented 10 degrees to the left.

You can test the validity of this technique yourself using the stimuli in **Figure 3.26**. Start by moving your eyes back and forth along the fixation line between the two gratings in Figure 3.26*a*. After about a minute of this adaptation, look quickly to the right at the fixation point between the vertical stripes in Figure 3.26*b*. The stripes in Figure 3.26*b* should now appear to be tilted slightly from their true vertical orientation (clockwise in the upper panel and counterclockwise in the lower—opposite the orientations in Figure 3.26*a*), just as predicted by the model of the human visual system based on the cat research and diagrammed in Figure 3.25. This **tilt aftereffect** strongly supports the idea that the human visual system contains individual neurons selective for different orientations.



**FIGURE 3.26** Stimuli for demonstrating selective adaptation. See the text for details. (After Frisby, 1980.)

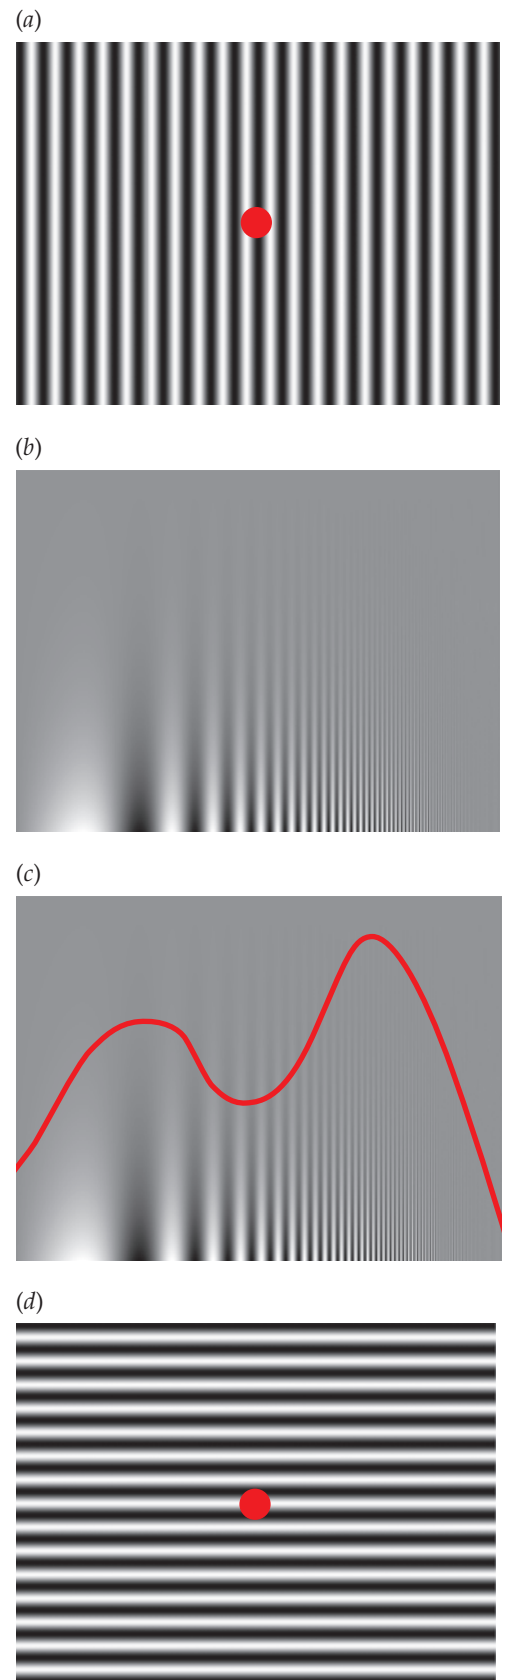
## Selective Adaptation for Spatial Frequencies


Selective adaptation also provides evidence that the human visual system contains neurons selective for spatial frequency. You can check this with the gratings shown in **Figure 3.27**. First look at Figure 3.27*b*, and make a mental note of your contrast sensitivity function (i.e., the inverted U-shaped area where the gratings fade into the gray background). Now adapt for about 10–20 seconds or so to the grating in Figure 3.27*a* and then quickly shift your gaze back to Figure 3.27*b*, and make a mental note of your contrast sensitivity function now. After repeating this procedure a few times, the outline of your CSF in Figure 3.27*b* should look something like the red curve in Figure 3.27*c*. It should have a notch (indicating reduced contrast sensitivity) for spatial frequencies that are close to the adapting spatial frequency (Figure 3.27*a*). This demonstration shows that adaptation to the high-contrast top panel is selective—resulting in a loss of sensitivity for spatial frequencies close to the adapting frequency, but not for spatial frequencies that are much higher or lower than the adapting frequency.

As noted earlier, selective adaptation causes the neurons most sensitive to the adapting stimulus to become fatigued. In this demonstration, neurons sensitive to the spatial frequency of the adapting stimulus have their contrast sensitivity reduced. That is, higher contrast is needed after adaptation for a test grating to be able to stimulate these neurons. Neurons responsive to much higher or much lower spatial frequencies are not fatigued by the adaptation procedure, so contrast sensitivity for these spatial frequencies is not affected.

**Figure 3.28*a*** shows more precisely how selective adaptation to a 7-cycle/degree grating produces a selective loss of contrast sensitivity at spatial frequencies of about 7 cycles/degree, with little or no loss at, for example, 1 cycle/degree or at 15 cycles/degree. After adaptation, the contrast sensitivity function has a “notch,” as if the detectors sensitive to spatial frequencies near 7 cycles/degree were selectively destroyed (luckily the effects of spatial-frequency adaptation are reversible!). From these measurements of contrast sensitivity before and after adaptation, we can construct the spatial-frequency tuning function shown in **Figure 3.28*b***. This function represents the change in contrast threshold (contrast threshold after adaptation, divided by contrast threshold before adaptation) plotted against the spatial frequency of the test grating. This curve represents the spatial-frequency tuning function for a “channel” that is most sensitive to a grating of 7 cycles/degree. The shape and selectivity of this channel are very similar to the spatial-frequency tuning functions for striate cortex neurons of cats and monkeys. **Figure 3.28*c*** illustrates that the contrast sensitivity function (gray curve) represents the “upper envelope” of the sensitivities of many spatial-frequency channels, each tuned to a different spatial frequency.

**FIGURE 3.27** A demonstration of adaptation that is specific to spatial frequency (SF). (a) The adapting grating. (b) A grating modulated in contrast (vertically) and spatial frequency (horizontally). This pattern allows the reader to visualize their own contrast sensitivity function (CSF; see also Figure 3.8). Before adaptation it should have the appearance of an inverted U. After adapting to (a), your CSF should look something like the red curve in (c). The red curve illustrates the effect of adaptation. The notch indicates reduced contrast sensitivity for spatial frequencies that are close to the adapting spatial frequency. (b from Robson and Campbell, 1997; courtesy of Izumi Ohzawa.)





**FIGURE 3.28** Spatial-frequency adaptation. (a) Selective adaptation to a 7-cycle/degree grating produces a selective loss of contrast sensitivity at spatial frequencies of about 7 cycles/degree, leaving a notch in the contrast sensitivity function. (b) Threshold elevation (the change in contrast threshold) following adaptation. (c) Physiologically measured spatial-frequency tuning functions for striate cortex neurons in monkeys (colored curves represent different neurons). (b after C. Blakemore and Campbell, 1969; c after Maffei and Fiorentini, 1973.)

### The Site of Selective Adaptation Effects

The adaptation experiments replicated here provide strong evidence that orientation and spatial frequency are coded by neurons somewhere in the human visual system. In cats and monkeys, we know that these neurons are located in striate cortex, not in the retina or LGN. Can we localize the orientation-selective and spatial frequency-selective neurons in humans?

As it turns out, we can do just that with a clever variation on the adaptation experiments. Repeat the orientation (see Figure 3.26) and spatial-frequency (see Figure 3.27) adaptation demonstrations, but this time view the adapting stimuli with your left eye only (keep your right eye closed during the adaptation period), then view the test stimuli with your right eye (close your left eye and open your right eye as you shift your gaze to the test stimuli). You should find that the tilt aftereffect and the decreased contrast sensitivity transfer from one eye to another, although the effect may be somewhat less pronounced than when you did the demonstrations with both eyes open. This transfer of adaptation from the adapted to the nonadapted eye is known as “interocular transfer.”

Now recall that information from the two eyes is kept completely separate in the retinas and in the two LGNs; no single neuron receives input from both eyes until the striate cortex. The transfer of adaptation effects from one eye to another thus implies that selective adaptation occurs in cortical neurons, just as we would predict from animal physiology studies.

### Combining Spatial Frequency and Orientation Selectivity

Before we leave the selective adaptation paradigm, let’s try one more demonstration. Adapt to the grating in Figure 3.27d for about a minute, then switch your gaze to the test grating in Figure 3.27b. You should find that the contrast sensitivity for this test stimulus is not reduced at all; the test grating should be equally visible both before and after adaptation. What does this result tell us about the coding of orientation and spatial frequency in human cortical neurons?


Remember that in Hubel and Wiesel’s studies with cats, each striate cortex neuron was selective for both a particular orientation and a particular spatial frequency. In other words, a cell responded to a test stimulus only if the stimulus was at (or at least near) the cell’s preferred orientation *and* close to the cell’s preferred spatial frequency. Our selective adaptation findings indicate that human cortical cells operate the same way. The test stimulus in Figure 3.27b has the same spatial frequency as the adapting stimulus, but a completely different orientation. Therefore, the two stimuli will stimulate

completely different neurons, and we should expect no contrast sensitivity reduction for the test stimulus.

### Spatial Frequency–Tuned Pattern Analyzers in Human Vision


Selective adaptation to spatial frequency, as well as other evidence, provides strong support for the notion, first suggested by Fergus Campbell and John Robson (1968), that the human contrast sensitivity function actually reflects the sensitivity of multiple individual pattern analyzers. These pattern analyzers are implemented by ensembles of cortical neurons, with each set of cells tuned to a limited range of spatial frequencies and orientations, and they are often referred to as **spatial-frequency channels**. Remember the initially unexplained falloff in the contrast sensitivity function at very low spatial frequencies? Although a number of different explanations have been suggested (e.g., lateral inhibition), the most likely explanation is that we simply have fewer neurons tuned to low spatial frequencies (De Valois, Albrecht, and Thorell, 1982).

The “multiple spatial frequency” model of vision implies that spatial frequencies that stimulate different pattern analyzers will be detected independently, even if the different frequencies are combined in the same image. Consider the compound grating pattern in **Figure 3.29**, made by adding a sine wave with frequency  $f$  to a sine wave with frequency  $3f$ . Graham and Nachmias (1971) found that the contrast sensitivity for this compound pattern was almost the same as the contrast sensitivity for detecting the individual components of the pattern separately. If the two component sine waves had stimulated a common pattern analyzer, then their effects on the analyzer should have been added, so contrast sensitivity should have been greatly improved.



**FIGURE 3.29** A compound grating pattern (right), made by the addition of a sine wave of frequency  $f$  (top left) to one of frequency  $3f$  (bottom left). (After Graham and Nachmias, 1971.)

**spatial-frequency channel** A pattern analyzer, implemented by an ensemble of cortical neurons, in which each set of neurons is tuned to a limited range of spatial frequencies.



**FIGURE 3.30** A complete image (a) and simulations of the high-frequency (b) and low-frequency (c) components of that image.

**FIGURE 3.31** Who is hidden behind the high-spatial-frequency mask in this image?



Why would the visual system use spatial-frequency filters to analyze images? One important reason may be that different spatial frequencies emphasize different types of information. **Figure 3.30** shows a face that has only high-frequency (panel *b*) or low-frequency (panel *c*) components of the face in panel (*a*). These images show that low frequencies emphasize the broad outlines of the face and high frequencies carry information about fine details. If we want to know how many people are in a scene, it is most efficient to consult our low-frequency channels. But if we want to know about the fine details—for example, whether the people are frowning or smiling—we must rely on our high-frequency channels.

It is widely agreed that at near-threshold contrasts, pattern analyzers operating at different scales of analysis are independent. At high contrasts, however, pattern analyzers do interact. You can experience this for yourself in **Figure 3.31**, where the high spatial frequencies introduced by the small blocks mask the low spatial frequencies that convey an underlying portrait of a famous American. Squinting your eyes will blur the blocks, minimizing the effect of the mask to a point where the face you've seen countless times on the 5-dollar bill will probably show through.

### The Development of Spatial Vision

William James (1890) described the infant's world as "a blooming, buzzing confusion." However, studies over the past several decades have shown that the visual system is much more developed at birth than we used to think. One of the difficulties in assessing vision in infants is that we can't simply ask them what they see. Rather, we have to think of tricky ways to coax that information from them. The most widely used method for studying infant vision is based on an observation by Robert Fantz in the early 1960s. What Fantz noticed is that if infants are shown two scenes, they invariably stare at the more complex scene (the

scene with the most contours). So, if an infant is shown two patches, one containing stripes, and the other uniform gray, the infant will prefer to look at the stripes. Of course, an infant who couldn't see the stripes would be equally likely to stare at the gray patch as at the striped patch. Thus, preferential looking is one important method used by infant researchers (grown psychologists studying infant vision, not babies in lab coats) to learn what infants can see and respond to behaviorally (Figure 3.32a).

The success of preferential looking depends on the willingness of babies to stare at stimuli near threshold level. An alternative approach, used with considerable success in more recent years, is to measure visually evoked electrical potentials (VEPs) to visual stimuli by attaching electrodes to the scalp and measuring the changes in electrical activity that are elicited by the changing visual stimulus (Figure 3.32b and c). Using the VEP, we can measure an entire contrast sensitivity function in as little as 10 seconds in a non-verbal infant.

### Development of the Contrast Sensitivity Function


The emerging picture suggests that, with increasing age, peak contrast sensitivity increases and the CSF peak shifts toward higher spatial frequencies. Low spatial-frequency sensitivity develops much more rapidly than high

**FIGURE 3.32** Assessing vision in infants. (a) Forced-choice preferential-looking stimuli (left) and the experimental setup (right). (b) Visual evoked potential (VEP) setup. (c) Results of a sweep VEP experiment in which the spatial frequency of the stimulus is swept (continuously varied from low to high spatial frequency), illustrating the extrapolated acuity. This particular experiment was done at 80% contrast. (Part c after Norcia et al., 1990.)


(a) Stimulus card

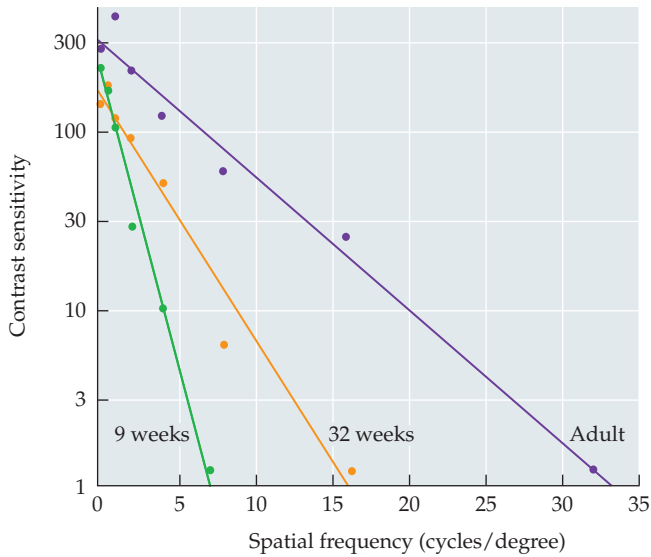


(b)



(c) Sweep VEP (grating acuity)





**FIGURE 3.33** The development of contrast sensitivity. (After Norcia et al., 1990.)

spatial-frequency sensitivity. Thus, peak contrast sensitivity may reach adult levels as early as about 9 weeks of age, whereas sensitivity at higher spatial frequencies continues to develop dramatically (Figure 3.33). There remains a roughly 20-fold difference in the contrast sensitivity of adults and 32-week-olds (Norcia et al., 1990).

What limits the development of acuity and contrast sensitivity? The primary postnatal changes in the retina concern differentiation of the macular region (Boothe, Dobson, and Teller, 1985). After birth, foveal receptor density and cone outer segment length both increase, as foveal cones become thinner and more elongated. There is a dramatic migration of ganglion cells and inner nuclear layers from the foveal region as the foveal pit develops during the first 4 months of life, and not until about 4 years of age is the fovea fully adultlike (Yuodelis and Hendrickson, 1986).

From birth to beyond 4 years of age, cone density increases in the central region, because of both migration of receptors and decreases in their dimensions. Both of these factors result in finer cone sampling (by decreasing the distance between neighboring cones). Alterations in cone spacing and the light-gathering properties of the cones during early development probably contribute a great deal toward the improvements in acuity and contrast sensitivity during the first months of life. The massive migration of retinal cells, and the alterations in the size of retina and eyeball (along with changes in interpupillary distance), may necessitate the plasticity of cortical connections early in life. Interestingly, the peripheral retina appears to develop much more rapidly than the fovea (Yuodelis and Hendrickson, 1986).

### The Girl Who Almost Couldn't See Stripes

Normal visual development requires normal visual experience. Abnormal early visual experience can have serious and often permanent consequences for seeing patterns, as illustrated by the story of a girl named Jane. Jane was born with a dense cataract (an opacity of the lens) in her left eye, which prevented clear patterns from forming on her left retina. In addition to causing form deprivation in the left eye, the cataract prevented Jane's two eyes from seeing the same images at the same time.

Studies in cats and monkeys dating back to Hubel and Wiesel in the early 1960s have shown that monocular form deprivation can cause massive changes in cortical physiology that result in a devastating and permanent loss of spatial vision (Wiesel, 1982). Hubel and Wiesel, and many other workers subsequently, demonstrated that there is a critical period of early visual development when normal binocular visual stimulation is required for normal cortical development. In cats and monkeys this critical period covers the first 3 to 4 months of life; in humans it is extended to something on the order of the first 3 to 8 years. During the critical period, cortical neurons are still being wired up to their inputs from the two eyes. This is a period of neural plasticity, when abnormal visual experience can alter the normal neural wiring process. If one eye is not receiving normal stimulation, the neurons that should be destined to respond to that eye do not become properly connected. In fact, some evidence suggests that these neurons are actually co-opted by inputs from the other, normally functioning eye.

If cataracts or other conditions, such as **strabismus**, in which one eye is turned so that it is receiving a view of the world from an abnormal angle, are left untreated during the critical period, the misplaced cortical connections can never be repaired. The result is often **amblyopia** (reduced visual acuity in one eye because of abnormal early visual experience) and an inability to perceive stereopsis (a lack of binocular depth perception; see Chapter 6). Correcting the condition later in life will thus have little effect, because the information from the now-functioning eye can never be properly conveyed to or processed by the cortex.

Luckily for Jane, her pediatrician found the cataract early, and the cataractous lens was surgically replaced by an artificial lens when she was 3 months old. The visual acuity in Jane's left eye just after the contact lens was inserted was 20/1200, about four times worse than the normal value for a 3-month-old. But when tested again 1 month later, acuity in her left eye had already begun to catch up with the acuity in her right eye. In fact, a recent study (Maurer et al., 1999) found significant acuity improvements only an hour after corrective measures had been taken.

## Summary

1. In this chapter we followed the path of image processing from the eyeball to the brain. Neurons in the cerebral cortex translate the array of stars perceived by retinal ganglion cells into the beginnings of forms and patterns. The primary visual cortex is organized into thousands of tiny computers, each responsible for determining the orientation, width, color, and other characteristics of the stripes in one small portion of the visual field. In Chapter 4 we will continue this story by seeing how other parts of the brain combine the outputs from these minicomputers to produce a coherent representation.
2. Perhaps the most important feature of image processing is the remarkable transformation of information from the circular receptive fields of retinal ganglion cells to the elongated receptive fields of the cortex.
3. Cortical neurons are highly selective along a number of dimensions, including stimulus orientation, size, direction of motion, and eye of origin.
4. Neurons with similar preferences are often arranged in columns in primary visual cortex.
5. Selective adaptation provides a powerful, noninvasive tool for learning about stimulus specificity in human vision.
6. The human visual cortex contains pattern analyzers that are specific to spatial frequency and orientation.
7. Normal visual development requires normal visual experience. Abnormal visual experience early in life can cause massive changes in cortical physiology that result in a devastating and permanent loss of spatial vision.

**strabismus** A misalignment of the two eyes such that a single object in space is imaged on the fovea of one eye, and on a nonfoveal area of the other (turned) eye.

**amblyopia** A developmental disorder that is characterized by reduced spatial vision in an otherwise healthy eye, even with proper correction for refractive error. Often referred to as "lazy eye."

Refer to the  
**Sensation and Perception**  
 website  
 ([www.sinauer.com/wolfe2e](http://www.sinauer.com/wolfe2e))  
 for activities, essays, study  
 questions, and other study aids.