Binocular Vision and the Perception of Depth

CHAPTER 8

8.1 INTRODUCTION

Up to now we've been concerned with monocular* optics; the cameras we considered had one (possibly compound) lens, the properties of vision we described required only one eye, and the optical instruments we discussed were mostly designed for the use of one eye. We now want to acknowledge that we, like many other animals, have two eyes and ask what might be the advantage of this ***binocular***† arrangement.

One advantage of the second eye is that it provides us with an increased *field of view.* Close one eye and you immediately notice that part of the scene previously visible is no longer in your field of view. Many animals, such as fish or rabbits, have eyes set in opposite sides of their heads, each eye providing a separate view of the world. Together, two such eyes, with practically no overlap in their fields of view, enable the animal to see a sweeping panorama. (The field of view of a rabbit is 360°, allowing it to see all around without turning its head, with only 24° overlap between the two eyes.)

However, your two eyes, placed in the front of your head, have fields of view that overlap considerably. (Your field of view is 208°, with 130° overlap.) Your eyes thus provide slightly different views of almost the same scene. Close one eye and hold up your thumb at arm's length. Now open that eye and close the other. Notice that your thumb appears to move against the background as you alternate eyes—each eye sees a slightly different view.

Nevertheless, normally you see only one view—somehow your brain combines the two images into one view of the world. In Figure 7.3, we showed the neural connections that allow this mixing of the signals from your two eyes. Notice that a given side of the brain gets signals from the corresponding points of each eye, allowing it to combine them into one view of the world. (There are occasions, however, when your brain isn't able to provide a single, smoothly combined, view, and you become aware of the two images produced by your two eyes. Hold up a finger in your field of view, fairly close to your eyes, and focus on a more distant object. You will notice two somewhat transparent fingers in front of the distant object. Alternately, focus on the finger, and you can see two images of the distant object. The TRY IT offers another example.)

Why bother with two images of the same scene? These slightly different images enable you to gauge the ***depth*** of a *three-dimensional* scene. Among animals, predators (such as cats) have their two eyes in front, with overlapping fields of view, to enable them to judge accurately the distance to their prey. On the other hand, animals (such as

FIGURE 8.1

Edward Hicks' "The Peaceable Kingdom" shows both predators and prey. They are easily distinguishable by the location of their eyes.

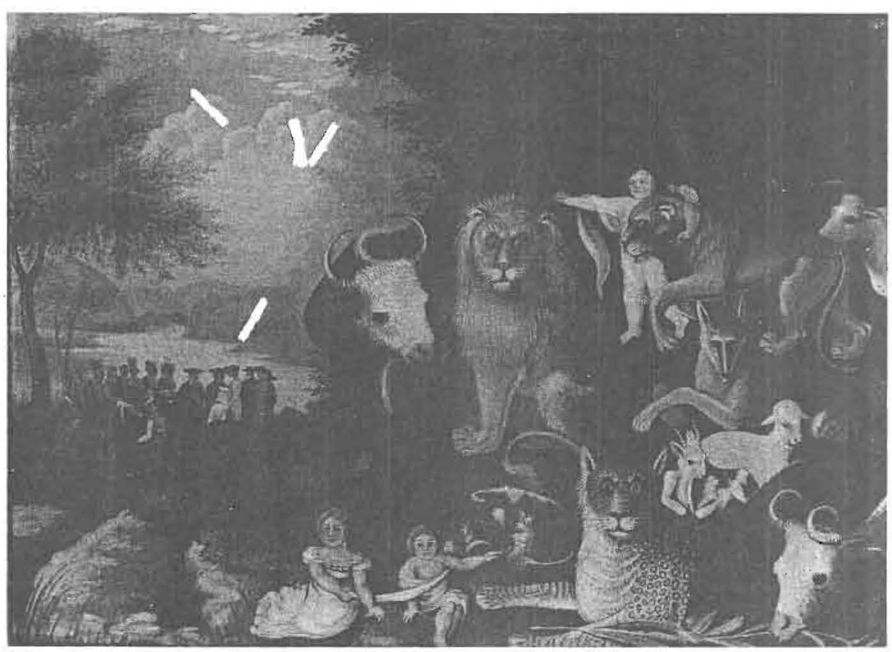

*Greek *monos*, single, plus Latin *oculus*, eye.

†Latin *bini*, a pair.

rabbits or deer) who are likely to be someone else's dinner have non-overlapping fields of view to give them the wide angle of view best for detecting predators (Fig. 8.1). Similarly, animals that leap about the branches of a tree (such as squirrels and our simian ancestors) must be able to gauge the depth of those branches, and correspondingly have two eyes in the front.

The brain, as we'll see, uses many ***cues*** in its determination of depth. Some require signals from two eyes, but others do not. These latter include the cues that artists rely on when they convey a feeling of depth in two-dimensional pictures. It is also possible to play off one cue against the other—to create scenes that provide conflicting visual depth cues. By examining the resulting illusions, we can learn about the way the brain processes the various cues in arriving at its depth determinations.

Let's examine a number of the techniques by which we visually fathom the depths of the world around us. We'll begin by separating the depth cues that can be used in a painting from those normally unavailable to the artist. Imagine that while you sleep an artist or photographer has made an extremely realistic picture of the view from your bedroom window and pasted it to the outside of your window so when you awake you see the picture (Fig. 8.2). How can you tell whether you are looking at such a painting or at the actual scene outside your window? In the next few sections, we'll discuss several techniques for distinguishing the two alternatives.

FIGURE 8.2

René Magritte, "The Human Condition I." How can we visually distinguish between the artist's rendering of the outside world and the world itself?

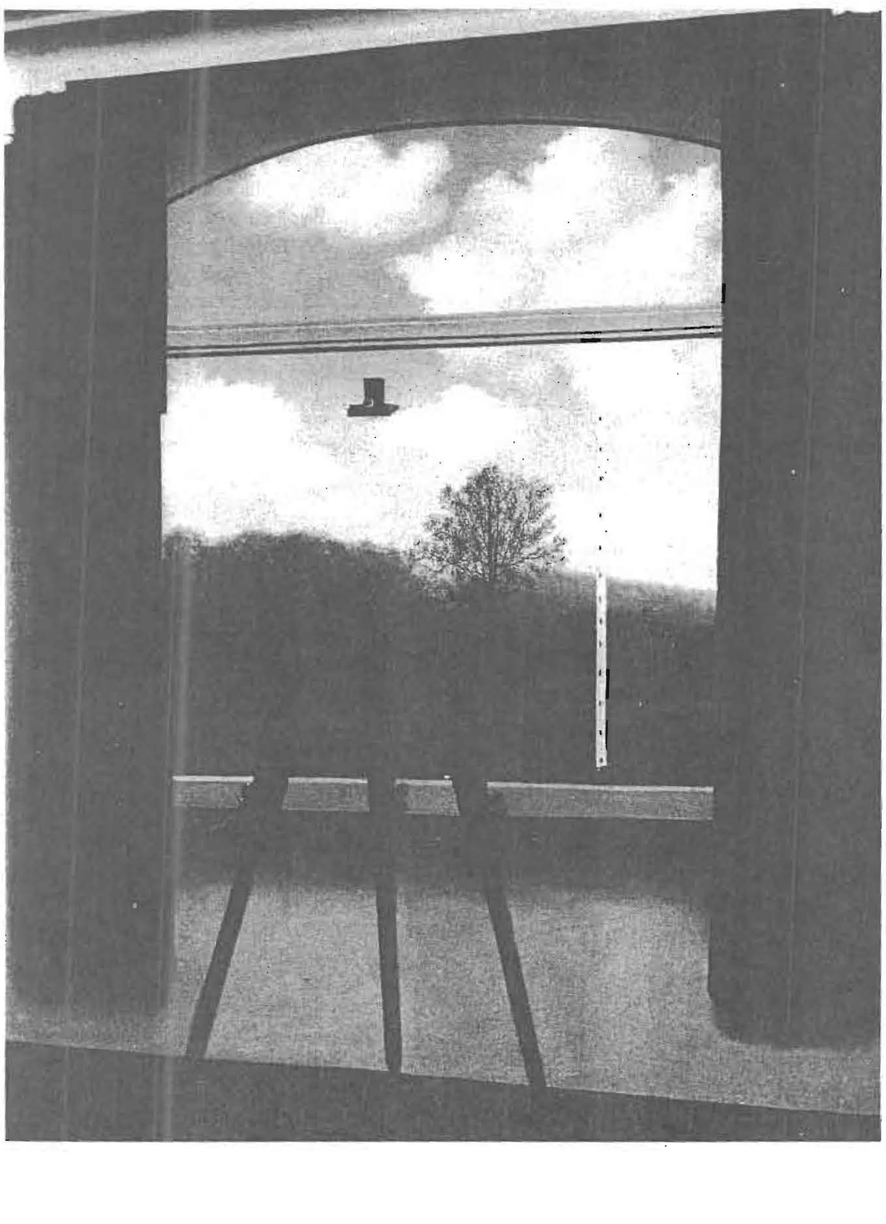

TRY IT

FOR SECTION 8.1
Two eyes provide two views

Hold one end of a string against your upper lip and pull the other end straight in front of you. You will see not one but two strings stretching out in front of you and crossing. These correspond to the images from your two eyes, as you can readily confirm by alternately closing each eye. The point at which they cross is the point on the string at which you "aim" your eyes (the point toward which your eyes converge). Try looking at different points of the string, beginning close to you and moving away, and notice how the cross-over point moves away from you as you do this. (Having a friend slide a finger along the string may enhance this effect.)

8.2 ACCOMMODATION

Just as you can measure the distance to an object by focusing your camera on it and noting the lens' position, the amount of *accommodation* necessary to focus your eye on an object tells you the object's distance. If you see an object clearly while your eyes are relaxed, you know that it is far from you. If, however, you must tense your ciliary muscles to make the object come into focus, then the object must be closer. Thus when looking out your window at the actual street scene, you would accommodate differently for objects at different distances in the scene. The artist, attempting to simulate that scene with a picture, must decide what is in focus and what is not. Once the artist makes that decision, no amount of accommodation on your part will sharpen the focus of an object painted out of focus.

If you were a chameleon, you'd rely very heavily on the technique of accommodation to gauge the distance to flying insects as you flicked them with your tongue. If you cover one eye of a chameleon, he maintains his high degree of accuracy in fly-flicking—binocular vision is not important here. However, give a chameleon a pair of glasses that change the amount of accommodation necessary and his tongue flaps futilely at the fleeing fly.

Humans, on the other hand, make little use of this technique, possibly because our potential meals are generally more than a tongue's distance away. Accommodation as a way of determining depth is at best only useful for close objects. If our mischievous artist confines herself to a distant scene and if you cannot get too close to the window, then both the painting and the actual scene would be in focus to your relaxed eye and you wouldn't be able to distinguish the one from the other by accommodation.

8.3 CONVERGENCE

If you look at a near object, the angle between your two eyes' directions of gaze is bigger than if you look at a distant object (Fig. 8.3). This angle is called the angle of ***convergence*** of your eyes. (For an object at the normal near point, 25 cm, the angle of convergence takes its maximum value of about 15°. It is only 1° for an object about 4 m from your eyes.) If your brain keeps track of the convergence of your eyes, it can determine the distance to the object that your eyes are viewing by using the rangefinder principle (Sec. 4.2D). As you scan a painted picture, your convergence, like your accommodation, remains unchanged, because all the objects lie in the plane of the picture. However, when you view objects at different depths in an actual scene, your convergence changes.

Like accommodation, convergence is most effective for determining depth in nearby scenes, but you make relatively little use of it. That you make *some* use of convergence, however, can be seen from the TRY IT.

TRY IT

FOR SECTION 8.3
Convergence and depth

Look at a distant street lamp. Cross your eyes (say by looking at your finger, which you hold in front of and just below the light). Notice that the light appears closer (and smaller) than before—the more your eyes converge, the closer. This is true even though you see two images, but it may be easier if you position yourself so that only one eye can see the lamp.

8.4 PARALLAX

To gauge depth, you rely much more heavily on the fact that your view is different from different positions—the phenomenon of ***parallax.**** With one eye closed, hold your thumb a few centimeters from this page, so it blocks your view of the word "parallax." By moving your head, you can change your view sufficiently so that you can see the word. This is possible only because your thumb and the word are at different distances from your eye. So, your view of different objects changes as you move, according to their distance from you (Fig. 8.4).

No matter how carefully our artist simulates other depth cues, as long as she confines herself to a flat canvas, she cannot overcome parallax. You need only move your head and compare the relative positions of the distant scene and the window glazing bars to determine if she has tried to trick you. Similarly, the painted finger of Lord Kitchener (Fig. 8.5) always points at you, no matter where you move with respect to the picture, but an actual finger points in one direction—when you move out of that direction, it no longer points at you. Because your view of Lord Kitchener's finger doesn't change, your brain may interpret the image as if he were rotating as you walk by, so as always to point at you. This provides the recruiting poster with a rather personal touch. (The FOCUS ON X-Ray

*Greek *parallaxis*, change.

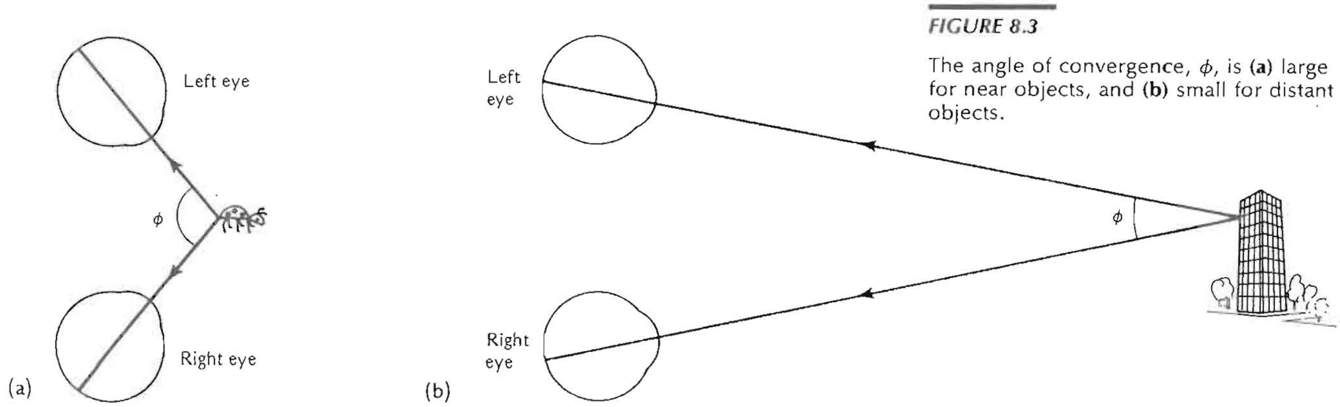

FIGURE 8.3

The angle of convergence, ϕ, is **(a)** large for near objects, and **(b)** small for distant objects.

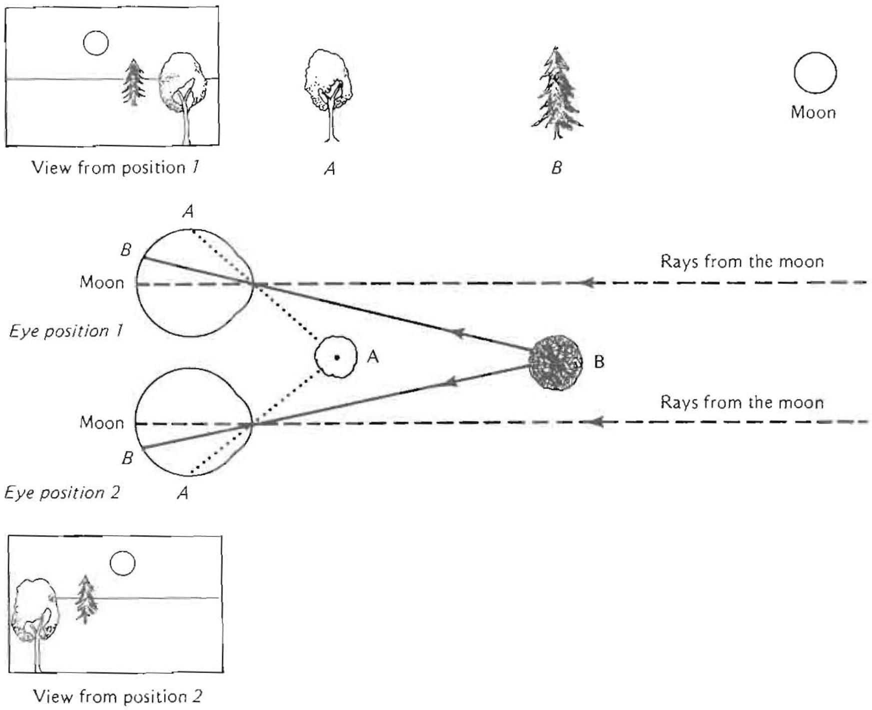

FIGURE 8.4

The locations of the images on the retina change as the eye moves from position 1 to position 2. The location of the retinal image of the nearest object, A, changes the most—from one extreme to the other. The image of a more distant object, B, moves less, while the parallel rays of the very distant moon are always imaged on the same place on the retina. This is why the moon seems to follow you as you look at it, say, from the window of a moving car. The closer the object is to you, the more it appears to move in the direction opposite to your motion.

FIGURE 8.5

Alfred Leete's 1914 recruiting poster of Lord Kitchener, a British hero of the Boer War.

Tomography gives an application of parallax.)

8.5
BINOCULAR DISPARITY

Those of us blessed with two eyes need not move in order to gain the benefits of parallax for gauging depth. The two eyes, separated by about $6\frac{1}{2}$ cm and with significantly overlapping fields of view, see slightly different views of any object they look at. This difference between the views of the two eyes ***(binocular disparity)*** thus provides a way of determining the distance of the object in sight. Your brain attempts to reconcile the two views by ascribing the difference to depth. You "see" one three-dimensional view of the world, rather than two two-dimensional views. If the images are too diverse, your brain cannot ***fuse*** them, and you get double vision. This happens sometimes when you're drunk or injured and your eyes don't properly converge to a single object. (As Shakespeare put it, "Methinks I see these things with parted eye, when everything seems double.")

Consider the two eyes viewing the cube in Figure 8.6. The views seen by the two eyes, shown in the figure, are slightly different. A common feature, say the front edge *aA*, is imaged at corresponding points of the two retinas. Another feature, say the left edge *dD*, is imaged at locations on the two retinas that do *not* correspond to each other. Some cells in the visual cortex of the brain respond strongest when a common feature occurs at corresponding points of the retinas. Other cells have their strongest responses at particular *differences* between the locations of the two

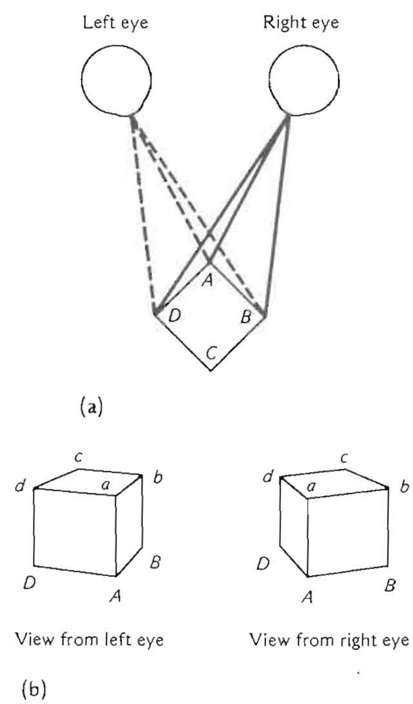

FIGURE 8.6

The two eyes looking at a cube see slightly different views of the cube. (a) View of eyes and cube. (b) Views seen by each eye. Note that the edge *da* subtends a smaller angle for the *right* eye than does *ab*, and consequently that eye sees *da* as shorter than *ab*. For the left eye, the situation is reversed.

retinal stimulations. Thus, one set of such cells in the visual cortex responds strongest to aA, while a different set responds strongest to dD. This difference in response leads to the viewer's perception that these features lie at different depths. (See the TRY IT's.) Under some circumstances, two objects will be seen in depth if the separations of their images on the two retinas differ by as little as 1 μm—less than the diameter of a cone!

Of course, binocular disparity is useful for depth determination only if your two eyes see different images. A horizontal clothes line looks the same to both eyes, so when you want to know where it is, it helps to tilt your head in order to introduce some disparity. (Objects that repeat themselves, such as bars on a cage, can offer confusing binocular disparity. If your eyes fuse the wrong bars together, the cage appears at a different distance than it should.) Distant scenes also present essentially the same view to your two eyes, so, like the other techniques discussed so far, binocular disparity is of little use for such scenes. But for relatively close objects it is an extremely effective way of gauging depth. In fact, by presenting two different views to your two eyes, it is possible to fool your brain into believing that there is depth even when there actually isn't. A variety of optical instruments, devices, and toys are based on this idea.

First TRY IT

FOR SECTION 8.5
Depth and chromatic aberration

The lenses of your eyes have some chromatic aberration—they bend blue light more than red. If you cover half your pupil, the resultant half lens acts like a prism, shifting the blue retinal image slightly away from the red image. By arranging that this shift is in opposite directions for the two eyes, you can produce binocular disparity—color differences can therefore be translated into differences in apparent depth.

In order to eliminate other depth cues as much as possible, you should look at blue and red patches separated and surrounded by black or white. Blue and red squares on one side of a Rubik's Cube can be used, as can a picture of an American flag on a white background. With pieces of cardboard or stiff paper, cover the outer half of each pupil while you look at the colored patches. Notice which color appears closer, then suddenly remove the cardboard pieces and notice the change. Now cut a piece of cardboard so it fits between your pupils, covering the inner half of each. (If you cut one of the vertical sides of the cardboard at an angle, you can get a good fit without precision measurements by raising or lowering the cardboard until it just doesn't block your view.) Look at the patches again, and notice that the depth has reversed from the previous case.

PONDER

Why did you need to use the cards?

Considering your half lenses as prisms, draw ray diagrams to convince yourself that the dispersion of the "prisms" is responsible for the effects you saw.

Second TRY IT

FOR SECTION 8.5
Increase your binocular disparity

The amount of depth perceived depends on the binocular disparity between your two views. The wider the separation of the points of view, the greater the apparent depth. Thus, stereoscopic aerial photographs are often taken from the two wing tips of the airplane. If you find the world too shallow, you can use the periscope described in the TRY IT for Section 2.4C to increase your binocular disparity. Hold the periscope horizontally and look through it with one eye while also looking with the other, unaided eye at the same object. A short periscope works best; about 6 cm will double your eye separation. The difference in apparent depth is most obvious if, after looking through it, you quickly remove the periscope while continuing to look at the scene.

Just as your brain interprets nerve signals from both of your eyes to produce the sensation of depth under normal viewing conditions, it can interpret signals accompanying afterimages with similar results. The persistence of positive afterimages also permits you to achieve an enhanced depth. Use the technique described in the first TRY IT for Section 7.7A, but here, since each eye is separately exposed to the window scene, care must be taken that the eyes are pointed in the same direction. After covering both eyes with your hands for 30 seconds, expose your right eye. To be sure that it is aimed properly, momentarily squint and point your eye at a distant point, such as the top of a distant tree. Then open your right eye wide for three seconds, keeping it pointed at the treetop. Next close and cover your right eye. Immediately move your head a few centimeters to the left and similarly expose your left eye to the scene, pointing it at the same treetop. Expose this eye for only two seconds. Close and cover the left eye (the right should already be closed and covered) and watch the stereo afterimage develop. Compare the depth in the afterimage to the actual depth.

To appreciate fully this "superstereo," compare it to reduced stereo, achieved by moving your head to the right between the exposure of your eyes in the order described above. The reduced effective separation between the eyes gives a flatter view. Moving your head about 13 cm to the right may result in a pseudoscopic view (Secs. 8.5A and 8.6H).

A. The stereoscope and related optical instruments and toys

How can we present two different pictures to the two eyes, so that each eye sees only the image intended for it? A simple technique is that used in the nineteenth-century **stereoscope.*** Here two photographs taken from slightly different angles are placed side by side. You view them through special lens-prism combinations, one for each eye (Fig. 8.7). The lenses produce distant virtual images of the photographs, while the prisms cause these virtual images to appear at the same place. Thus, the lenses assure the proper accommodation

*Greek *stereos*, solid.

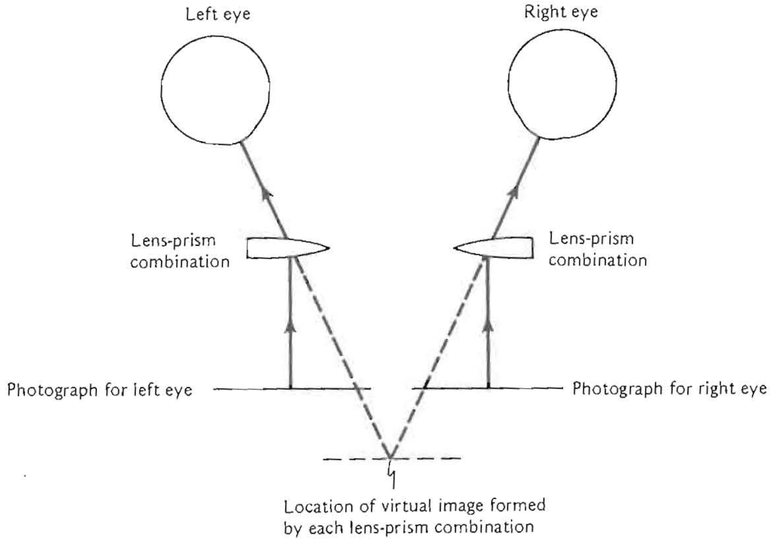

(a)

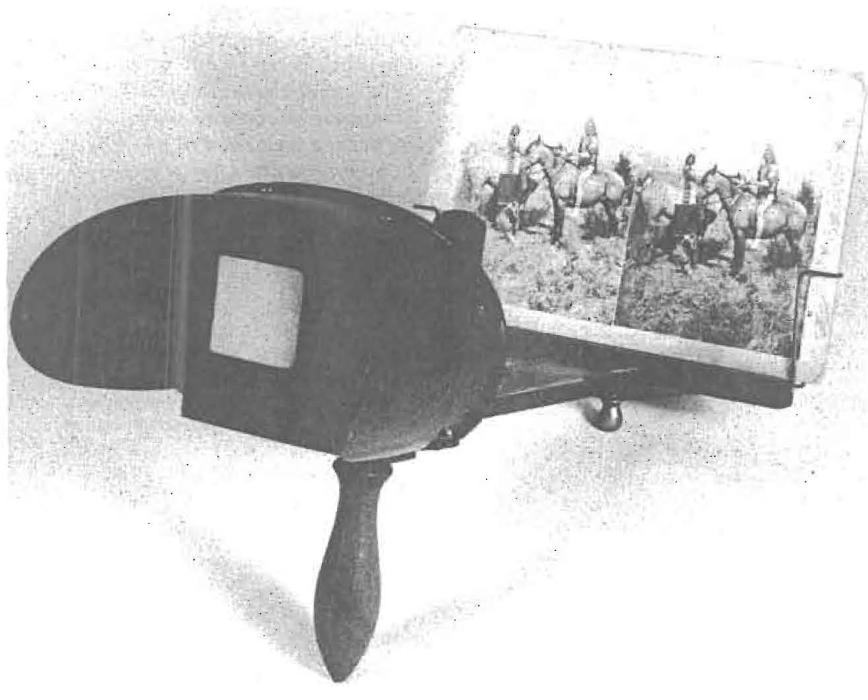

(b)

FIGURE 8.7

(a) The principle of the stereoscope. The amount of depth perceived depends on the disparity between the two photographs. **(b)** A stereoscope from the 1880s viewing a stereo pair of photographs of Indians.

and the prisms the proper convergence for an object at one, average distance. But the binocular disparity resulting from differences in the photographs often produces quite a realistic impression of depth. (The modern "Viewmaster," used for scenic views or as a toy, is just an inexpensive version of Hermann von Helmholtz's 1866 stereoscope.)

With a little effort, you can obtain depth from stereoscopic pairs of pictures without paraphernalia. (That is, most of you can. About 2% of the population lack the stereoscopic vision to see depth from stereo pairs, even if they have two good eyes.) Figure 8.8 shows a stereoscopic pair. If you hold a piece of cardboard between your eyes so it blocks each eye's view of the other eye's photograph (being careful to avoid shadows), you may be able to fuse the two pictures even though the convergence and accommodation is wrong. It may take a little time and concentration before the images fuse. If you are able to do this, you can then use this talent to pick out small differences in otherwise identical photographs, since only the different parts will appear in depth. This depth effect has been used to notice changes in the position of stars. Two photographs of the night sky, taken at different times, are viewed stereoscopically. Those stars that have moved appear in depth. (A better way is to view the two photographs in rapid succession, repeatedly, and then look for the stars that appear to move—Sec. 7.7A.)

Stereo pairs of pictures may be obtained in a variety of ways. Cameras have been made with side-by-side lenses that produce two photographs simultaneously. You can make stereo pairs with an ordinary camera by photographing a still scene, then moving the camera to one side, and taking another picture. Stereo pairs are produced on the scanning electron microscope by taking a picture, tipping the microscopic sample slightly, and taking another picture. You can even make drawings, such as Figure 8.6b. (See the first TRY IT.)

What happens if you interchange the two pictures, say of Figure 8.6b, so the right eye views the left eye's picture and vice versa? The left eye will then see edge *da* as shorter than *ab*, while the right eye sees the reverse. This is just what these eyes would see if they viewed the three-dimensional object shown in Figure 8.9. That is, instead of the original front edge *aA* appearing closer to the viewer, it appears *farther* than the side edges *dD* and *bB*. Such a view, where the parts of the original object that came forward

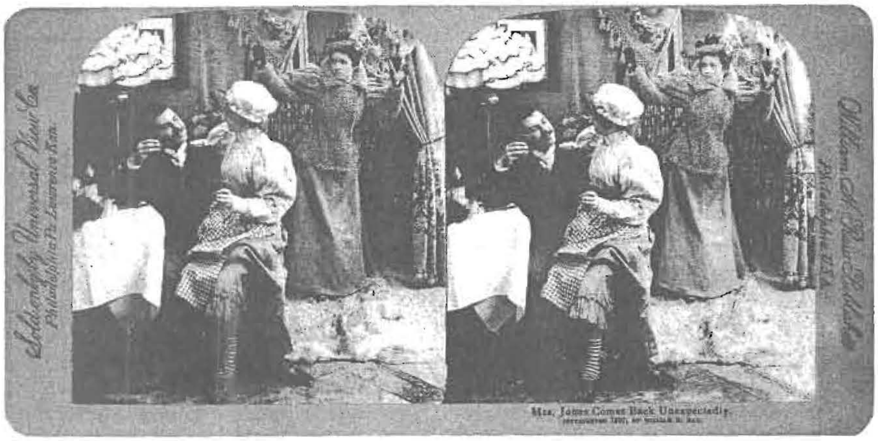

FIGURE 8.8

Stereoscopic photograph: "Mrs. Jones Comes Back Unexpectedly" (1897).

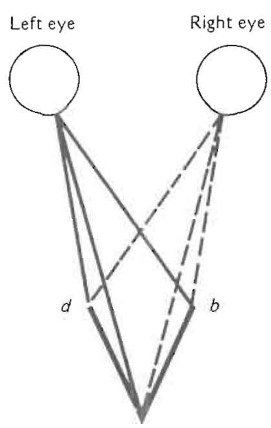

FIGURE 8.9

The two eyes looking at this object would have views corresponding to those shown in Figure 8.6b, but with the views interchanged. The edge *da* now subtends a smaller angle for the *left* eye, so that eye sees *da* as shorter than *ab*.

appear to go back and vice versa, is called ***pseudoscopic.****

Various other optical devices that offer separate images to each eye present binocular disparity and hence a three-dimensional view to the observer. (The second TRY IT gives an unusual example.) The most common such instrument is binoculars—actually binocular telescopes, one for each eye. Binocular microscopes are more complicated. Because the image produced by a microscope is usually *inverted* (Fig. 6.9), simply providing each eye with its own microscope directed at the same object would produce a pseudoscopic view.

PONDER

Why does each eye viewing an inverted image result in a pseudoscopic view? A lens inverts the image (turns it upside down) by reversing both up-down and *left-right. Draw these reversals separately for the two views of Figure 8.6b, first a left-right reversal and then an up-down reversal of* that. *A mirror may help in each case. Which reversal causes the stereo pair to result in a pseudoscopic image?*

For this reason, some binocular microscopes do *not* offer a stereo view. They simply have a binocular eyepiece so that you can look at the *same* image with both eyes. Other binocular microscopes do offer a proper stereo view by inserting an *erecting system* in each of two microscopes. This may be a set of Porro prisms (Fig. 2.56c), as in binoculars, or some other system involving a reflection.

Another way of providing separate images to the two eyes from two-dimensional pictures does not require that the two pictures be separated. The pictures are superimposed, and the viewer is given special eyeglasses that allow each eye to see only the image meant for it. This may be done in two ways: using *color* or using *polarized light* (Sec. 1.3B and Chapter 13). Consider first the use of color. Here the two eyes' views are printed in different colors, say red and green, as in Plate 8.1. (Such a picture is called an ***anaglyph.****) The viewer is given a different colored filter for each eye. If the filter is red, the green printing appears black, while the red printing is almost invisible. If the filter is green, the reverse is true. Thus, the eye with the red filter sees only the green picture, while the eye with the green filter sees only the red picture. Despite the difference in colors, the images can usually be fused into a three-dimensional view. (To produce three-dimensional shadows this way, see the third TRY IT.) This technique was used in early "3-D" movies, but because the color is used for the stereo effect the resultant three-dimensional view is not in full color. Full-colored 3-D movies can be achieved by using polarized light. The idea here is to project the two images with light of different polarization, say one with vertically polarized light and one with horizontally polarized light. The viewer looks through polarizing filters. The filter over one eye allows only the vertically polarized light to pass, while the filter over the other eye, oriented perpendicularly to the first, passes only the horizontally polarized light—but each "in living color." (Actually, the polarization is usually at 45°, one running between upper right and lower left, and the other between upper left and lower right.)

It is even possible to make stereo pairs of pictures where each member of the pair, by itself, doesn't

*The name (Greek *pseudes*, false) given by Sir Charles Wheatstone who, in 1838, first constructed a stereoscope. The first TRY IT shows how to reproduce his model.

*Greek *ana*, upward, plus *gluphein*, carve. The term was first used for low relief carvings.

FIGURE 8.10

The eye and brain can pick out the dog in this seemingly random array of dots.

look like anything recognizable, but when the pair is properly viewed, a three-dimensional figure becomes apparent. This can be done with patterns of random dots. The eye and brain are quite good at finding order and sorting out familiar patterns when presented with a seemingly random array of dots (Fig. 8.10). (See the fourth TRY IT for other examples.) If we present to the two eyes arrays of dots that, viewed monocularly, have no pattern but do have a definite stereoscopic relation with one another, the arrays will exhibit a pattern of depth when viewed properly binocularly. We can make such arrays by the method shown in Figure 8.11a. To construct the left eye's view, we need only make a rectangle *(R)* covered with some pattern. To construct the right eye's view, we simply cut out a square *(S)* from the center of the left eye's view, slide it over, and fill the gap produced *(Y)* with more pattern. Any pattern will do, including a random array of dots, providing there *is* some pattern that the brain can correlate between the two views. (A solid color would not work because we are not going to draw any boundaries on *S*.) With random dots, each eye's view shows no sign of the square *S* when viewed monocularly. Nevertheless, by comparing the two apparently random patterns, the brain can notice the correlation (some of the dots are moved with respect to the others) and attribute this correlation to depth (Fig. 8.11b and Plate 8.2). These random dot stereograms, developed by Bela Julesz, have been extremely useful in sorting out which states of visual processing occur before stereoscopic fusion and which after.

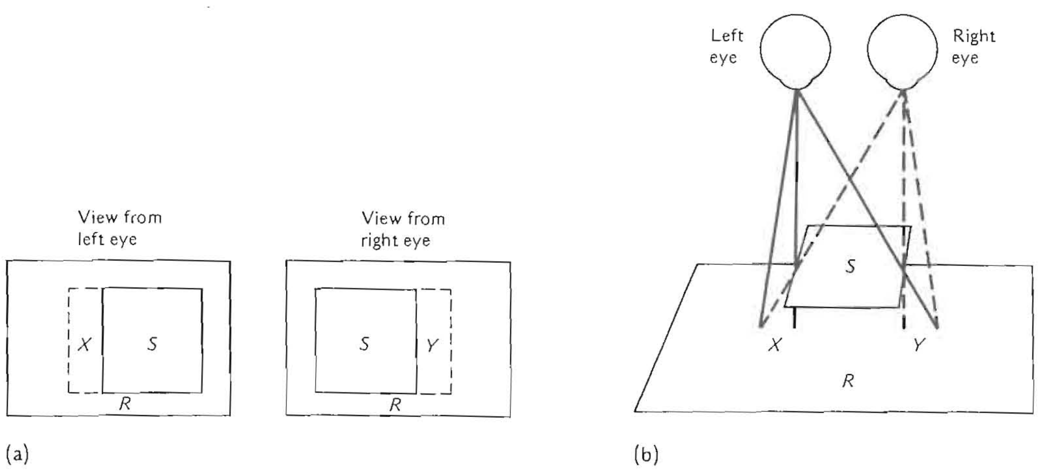

FIGURE 8.11

Constructing a stereogram. **(a)** The views of the two eyes. The left eye sees *X* but not *Y*. Everything else in the rectangle *R* is the same for the two views. Properly viewed stereoscopically, the two eyes will fuse these views into a three-dimensional image even if *R* and *S* are covered with random arrays of dots. **(b)** The three-dimensional view is that of a square *S* floating above the rectangle *R*. For the left eye, *S* obscures the region *Y* of *R*. For the right eye, *S* obscures the region *X* of *R*.

First TRY IT

FOR SECTION 8.5A

Constructing and viewing stereo pictures

Stereo pictures may be made and viewed in a variety of ways. For viewing with a stereoscope, you need two views from slightly separated positions. To do this photographically, choose a scene that has objects at a variety of distances from the camera. Get the greatest possible depth of field by using the largest f-number on your camera. It is best to rest your camera on a flat surface, so you can carefully slide it to one side for the second picture. The distance between camera positions depends on the way the pictures are to be viewed and whether you want normal or exaggerated depth. If the object is distant (beyond 2 m) and the pictures are to be viewed through magnifying glasses (as in most stereoscopes) so the image you see is distant, then you'll want to slide the camera sideways about $6\frac{1}{2}$ cm, the distance between your pupils. (The focal length of the magnifying glasses should equal that of the camera lens so you see the same perspective the camera did.) Try taking several pictures at various separations—the greater the separation the greater will be the apparent depth, but if the separation is too large, you won't be able to fuse the two pictures. Very distant shots, from airplanes or tall buildings, are usually taken with a separation between pictures of about $\frac{1}{50}$ the distance to the subject in order to increase the apparent depth. For close-up photographs, you should rotate the camera as you slide it, so it is always directed at exactly the same point on the object. For very close objects, you'll want to slide the camera less than $6\frac{1}{2}$ cm. (For magnifications greater than one, a good rule of thumb is to divide $6\frac{1}{2}$ cm by

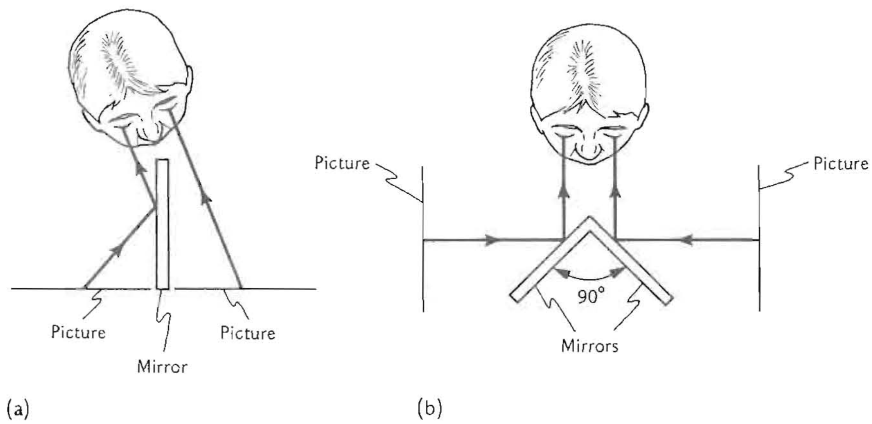

FIGURE 8.12

(a) Single-mirror stereoscope.
(b) Wheatstone's stereoscope.

the magnification.) Be sure to record which picture was taken from the left side and which from the right.

The stereo pairs may be viewed without a stereoscope by the method described in Section 8.5A. Fusing the images may be easier if the views are crossed—that is, the photograph for the left eye is placed on the right. (Looking at a pin held in front of you may assist you in crossing your eyes.) In another technique, you print one of the stereo pair of photographs left-right reversed, and instead of using a piece of cardboard between your eyes, you use a mirror (Fig. 8.12a). Alternatively, you can use Wheatstone's 1838 design to construct your own stereoscope. You will need only two small plane mirrors and some way of supporting them (chewing gum, if all else fails). Position the two mirrors at an angle of 90°, with the reflecting surfaces facing outward (Fig. 8.12b). Looking into the mirrors from the vertex of the angle, the left eye will see a view to the left and the right eye a view to the right. Position the photographs parallel to and facing each other and far enough from your eyes so that you can focus on them (25 cm, including the reflection). If you have two identical converging lenses, you can hold them in front of your eyes and place the photographs at the focal plane of each lens. Since there are mirrors (which reverse the images), you should print the pictures left-right reversed or you'll get a pseudoscopic view.

You can also draw stereo pairs for your stereoscope. This may be done in the same way as Figure 8.6 was drawn. If you plan to view the pictures at 25 cm, you should draw a diagram like Figure 8.6a, but full scale, with the object at 25 cm from the eye, and the eyes separated by $6\frac{1}{2}$ cm. To obtain the proper horizontal distances for your version of Figure 8.6b, draw a horizontal line through the object in your version of Figure 8.6a and measure the distances between the locations where rays cross that line. For example, the horizontal distance between points c and d for the left eye would be the distance, along that line, between the rays that go from those points to the left eye. Alternatively, if you have a sufficiently simple real object and some artistic talent, you can directly draw the pictures as seen by each of your eyes.

Very simple figures can be drawn without artistic talent. For example, two circles, one inside the other, can be drawn with a compass. Make the outer circle about 6 cm across and the inner circle 2 cm across. For one eye's view, the center of the inner circle should be a few millimeters off center to the right, and for the other eye's view, it should be off center the same distance to the left. The resultant stereo pair can be viewed in all the usual ways. The single-mirror technique of Figure 8.12a is particularly simple here, since for this case you draw the inner circle off center in the same direction for both eyes. (Why?)

Instead of a stereoscope, you can use colored filters and make anaglyphs. The colored filters can be the cheap plastic kind. You can make a photographic anaglyph by taking a double exposure on color film, holding first one filter in front of the lens, and then the other, and moving the camera to the side between exposures. Alternatively, you can draw an anaglyph by one of the above methods, using colored pens, pencils, or crayons, and drawing the two pictures one on top of the other. The colors with which you draw should be matched to the filters. You want colors that are easily seen through one filter but are invisible through the other.

You can also use polarized light (Chapter 13) to project stereo pictures. You'll need four seperate polarizing filters and two projectors, as well as a "metallic" projection screen that does not destroy the polarization. You can use polarizing sunglasses, but you'll either need four pairs, or you'll have to break them in half. (If you still have your glasses from a 3-D movie, you can get away with two pairs of them, only one of which you need break.) Take two stereo pictures, as above, to make color slides. Tape a polarizing filter over the lens of each projector and project the two slides, with their images the same size and overlapping as much as possible. View the screen with a polarizing filter over each eye. These should be oriented perpendicularly to each other (i.e., so you can't see through them when they are held one behind the other). Rotate the filters over the projectors until each eye can see only the image from the projector meant for that eye.

Second TRY IT

FOR SECTION 8.5A
The dark axle

You'll need a metal ring 6 cm across or larger (like those used in macrame). In a dim room, illuminate a table top at a low angle, say with a flashlight. Spin the ring as you might spin a coin—it will look like a transparent sphere spinning around a dark axle. The spherical appearance is due to persistence of your vision in the dim light; the dark axle comes about because the ring looks extra dark when its front hides its back. This occurs for each eye when the ring is lined up with that eye. Convince yourself that the resulting binocular disparity locates the axle in the center of the sphere.

Third TRY IT

FOR SECTION 8.5A
Three-dimensional shadows

For this you'll need two lamps that throw good sharp shadows, a red and a green light bulb, a white bed sheet, and a red and a green filter. (The filters should be colors that make the biggest difference when held in front of the two bulbs.)

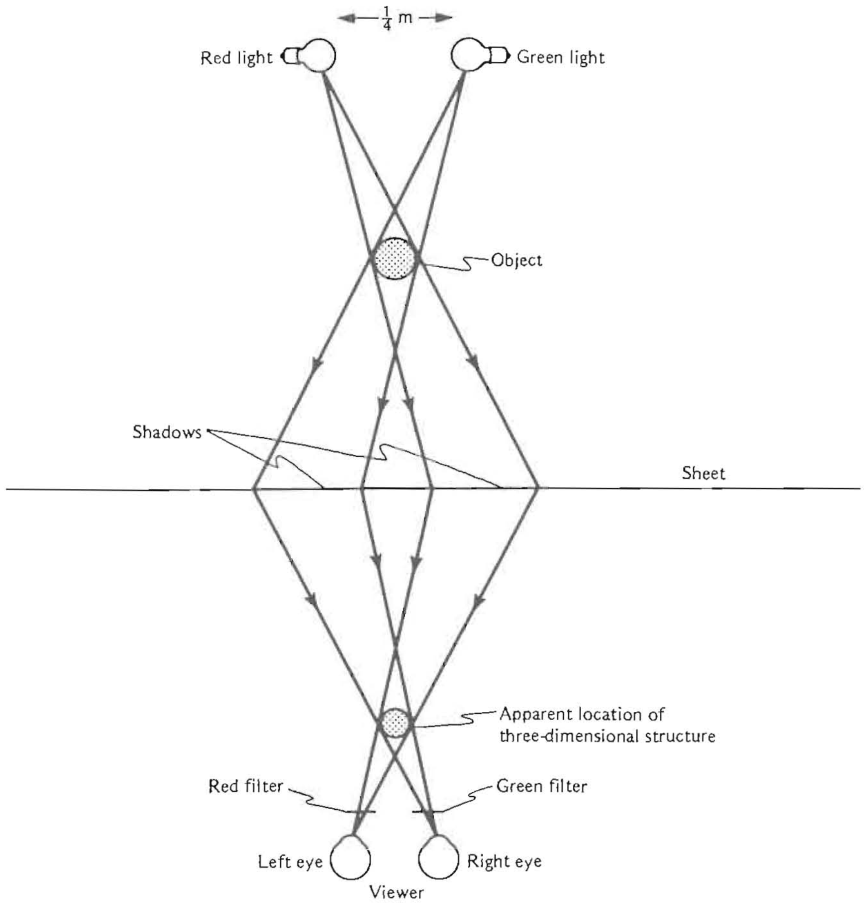

FIGURE 8.13

Set-up for three-dimensional shadows.

Arrange them as in Figure 8.13. With all other lights off, an object behind the screen will cast two shadows on the sheet, one due to each light. Viewed through the filters, the shadows will fuse into one three-dimensional shadow, apparently in front of the screen. A good object to use is a person. As he walks back from the screen toward the lamps, his shadow will appear to move forward from the screen toward the viewer.

Fourth TRY IT

FOR SECTION 8.5A
Order out of chaos

Excellent random arrays of dots can be easily found on your TV screen if you simply tune to a nonbroadcasting channel, so the screen is filled with "snow." This consists of randomly placed bright spots which, because the TV scans through all 525 lines in $\frac{1}{30}$ sec (in North America), change completely 30 times a second. Therefore, the dots seem to jump about at random, as the eye associates motion to the disappearance of a dot at one point and the subsequent appearance of another dot at a nearby location. Thus, the snow seems to consist of randomly placed dots randomly jumping about—a truly chaotic scene. It is, however, quite easy to "train" these dots to behave in a more orderly fashion, thanks to the ability of your brain to organize the visual information presented to it. All you need do is give your brain the appropriate hints.

Form a narrow channel between two pieces of paper held against the TV screen, separated by a few centimeters. Notice that the dots now seem to flow along the channel. Now make a loop out of wire or string, about 5 to 6 cm across, and hold it against the screen. Notice how the dots jump around within and without the loop, but do not seem to cross the loop. This remains true even if you move the loop across the screen—the dots seem lassoed by the loop and move along with it. There seems to be a void immediately behind the loop as it slides through the sea of jittering dots. You can form a furrow through the dots simply by moving your fingertip across the screen.

B. Lenticular screens

Another device that allows different images to be seen from a two-dimensional picture is the **lenticular screen.** Consider a sheet of cylindrical lenses (called **lenticules**) resting on a flat picture at the focal plane of the lenses (Fig. 8.14). (Such a sheet can easily and cheaply be made of plastic.) Light originating or scattered from any point in the picture plane is made parallel by the lenticules. The direction in which the light emerges from the lenticules depends on the location of the source in the picture plane. Light from a point slightly to the left of the axis of any of the lenses will emerge directed toward the right. Thus, light from all the points labeled R in the figure will emerge headed to the right, as shown. Similarly, light from the points labeled L will emerge directed to the left. This means that an observer located at O_R, looking slightly to the left, will see only the points R, while an observer at O_L, looking slightly to the right, will see only the points L.

If the picture plane consists of alternating strips of two different pictures, one at the points R and the other at the points L, then you see two different pictures, depending on the angle at which you view the lenticular screen. It may seem a bit tedious to cut a picture into little strips and paste them on the picture plane, but the strips can be made quite simply photographically by using the lenticular screen itself. Projecting a picture from O_R toward the points R will only expose the points R. Another picture can be projected from O_L toward the points L, only exposing those points. In fact, it is possible to get several more pictures between those

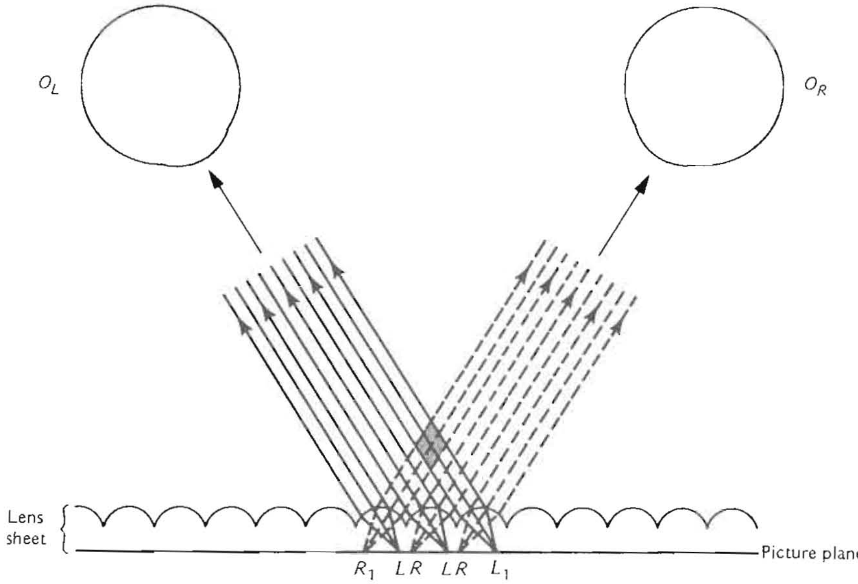

FIGURE 8.14

The lenticular screen. A sheet of cylindrical lenses is backed by a picture, lying at the focal plane of the lenses. Light originating (or scattered) from the points labeled R is made parallel by the lenses and sent to the upper right. Light from the points labeled L is made parallel and sent to the upper left. If we place two different pictures, intermeshed, at the points R and L, we see the R picture only if we look from O_R, the L picture only if we look from O_L.

points, each corresponding to a different direction of exposure, and later of viewing. The resultant device, which presents a different picture to you as you move past it, changing your angle of view, is sometimes seen in advertisements, picture postcards, and even political campaign buttons. It is reminiscent of the nineteenth-century puzzle pictures that were painted on a series of wedges (Fig. 8.15). The artist Yaacov Agam has made modern abstract (usually geometrical) pictures on such wedges, so that the viewer sees several different patterns as she moves about. He has also used lenticular screens similarly to present a series of different images to the viewer. (The latter he modestly calls "Agamographs.")

FIGURE 8.15

(a) By painting one picture on the R faces of the wedge and a different one on the L faces, an artist can present two pictures. Viewed from the right, the observer sees only the R picture; from the left, the observer sees the L picture. **(b)** Two views of such a Victorian puzzle picture. **(c)** Another puzzle picture that shows three views. Two views are drawn on opposite sides of the parallel slats.

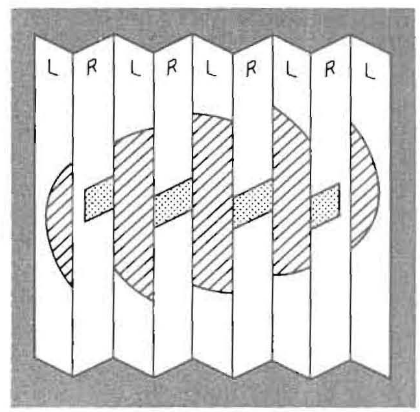

(a)

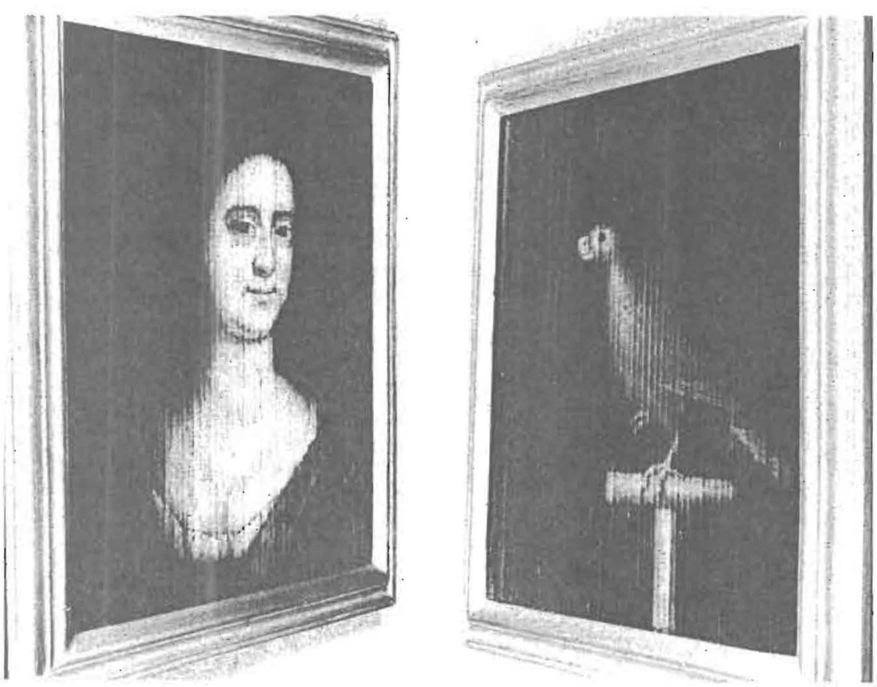

(b)

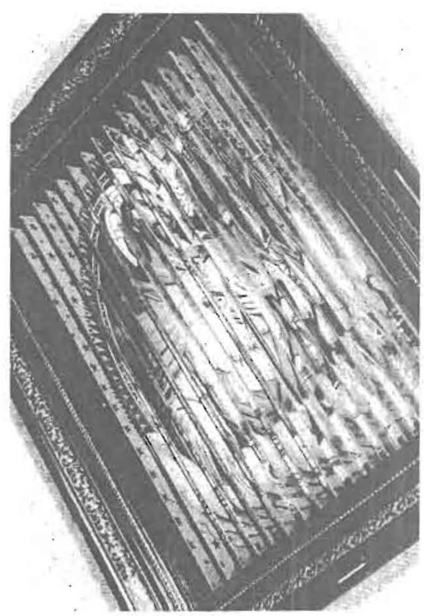

(c)

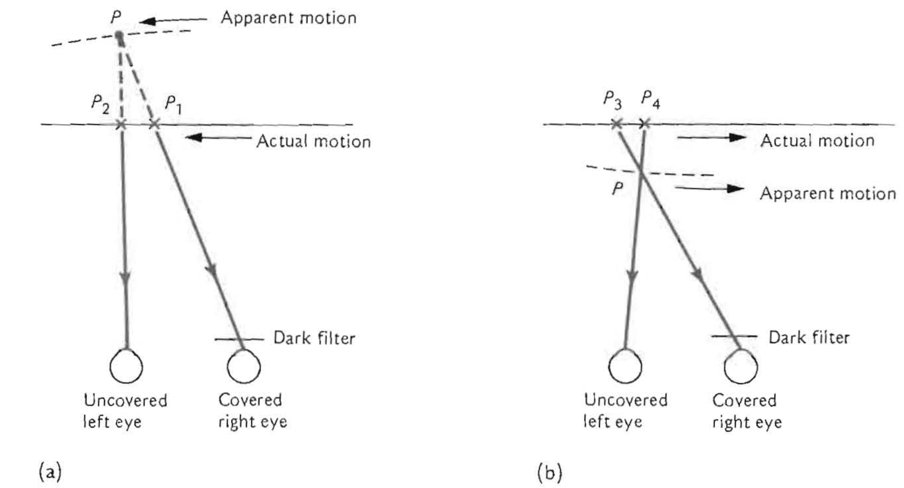

FIGURE 8.16

The Pulfrich pendulum. **(a)** The pendulum moves to the left. The covered right eye sends a delayed signal to the brain, so the viewer "sees" the pendulum at P_1 with that eye at the same time that she "sees" it at P_2 with the uncovered left eye. Her convergence and binocular disparity then tell her that the pendulum is at *P*, behind the actual plane of motion of the pendulum.
(b) The pendulum has now reversed its direction. Again, the right eye "sees" the pendulum at the delayed location, P_3, which is now to the left of the location at which the left eye "sees" the pendulum, P_4. The point *P*, on which both eyes agree, is now in front of the actual pendulum plane.

If the *R* pictures and *L* pictures correspond respectively to the right- and left-eye views of a three-dimensional object, then a viewer with his left eye at O_L and his right eye at O_R sees a three-dimensional view of the object. Each eye sees only the appropriate view necessary for the correct binocular disparity. The angle between the two views should not be too great, corresponding to the natural convergence the viewer would have when looking at the object. Such three-dimensional pictures can be found on picture postcards and sometimes on children's books. Stereo prints from the multiple pictures of the four-lens Nimslo camera are also made this way. It is important when viewing these pictures that the cylindrical lenticules should run vertically. If you rotate the picture so they run horizontally, both your eyes see the same image and the picture loses its depth. (However, if you then slowly tip the picture backward, the image you see will change.)

The picture may even appear in front of the lenticular screen. Suppose that the picture plane is painted a solid color, say red, except for two points (labled R_I and L_I in Fig. 8.14) that are painted green. If a viewer is positioned with his left eye at O_L and his right eye at O_R, both eyes see the same thing, red, except at the point where the rays from R_I and L_I intersect. This point, shaded in the figure, appears green. Since this green point lies in front of the picture, the viewer sees a green spot floating in front of a red background. This effect has been used by the artist Frank Bunts, whose dizzying pictures may have spots floating as much as a foot in front of the picture plane. It is well suited to producing three-dimensional geometrical patterns, and simplified versions are sometimes found in toys and novelties.

*C. The Pulfrich phenomenon

Another phenomenon that produces a stereoscopic effect by presenting different views to the two eyes relies on the fact that the latency time of the eye (Sec. 5.3C) varies with the intensity of light. In dimmer light, the retinal cells respond more slowly, as they must wait until there is sufficient exposure. Thus, an eye that is covered with a dark filter will respond more slowly to the same subject than one that is not.

Suppose you put a dark filter over one eye and look binocularly at a pendulum swinging in a plane perpendicular to your direction of vision (Fig. 8.16 and the TRY IT). Your brain gets the information about the location of the pendulum from the covered eye a little later than from the uncovered one. Thus, your two eyes locate the moving pendulum at two different positions (your covered eye "sees" the pendulum slightly in the past, compared to the uncovered eye). When the pendulum swings toward the side of your uncovered eye (Fig. 8.16a), the views of your two eyes, extended backward, converge behind the plane of the pendulum. The convergence and binocular disparity information lead you to "see" the pendulum behind its actual position. When the pendulum reverses its direction (Fig. 8.16b), its apparent path is in front of its actual path. Thus, the pendulum appears to swing out of its actual plane, in an ellipse.

This effect is named for Carl Pulfrich, who never saw it. He had been blind in one eye for sixteen years before the phenomenon was first noticed as a nuisance in an optical instrument that he designed. His colleague Ferdinand Fertsch actually suggested the explanation. The phenomenon was once used in diagnosing syphilis. Currently, it is used for the differential diagnosis of optic neuritis occuring in multiple sclerosis. In this case, the difference in response time of the two eyes is actual, so the patient sees the effect without any filter.