Chapter 8

The Third Dimension

Chapter Outline:

The Third Dimension

With each succeeding chapter, some new feature has been added to the visual world that increases complexity of the visual world that we can describe how the visual system operates. Up to this point the world has been flat. Yet we do not experience the world as flat. Even when the world really is flat such as in movies, realistic paintings, and photographs, we get a very powerful impression of depth. Take the photograph in Figure 8.1, which is a picture of one of my favorite places. You are looking northwest across Mount Desert Island in Maine back towards the mainland. It is easy to see what parts of the photograph represent objects that were closer or farther away than other objects. Moreover, the impression of depth is immediate and apparently effortless. There are trees in the foreground. Beyond those trees is a forest which goes up to the edge of a bay, Frenchman's bay. In the bay are island seen as farther from you than the island from which the photograph was taken. In the far distance you can make out hazy mountains. The ease with which a flat photograph can be seen to represent a world with depth suggests that the visual system is heavily wired to perceive a world that is three-dimensional.

The Problem of Depth

All of these observations serve to hide a rather interesting feature of depth perception. The retina is essentially flat. While it is curved along the back of the eye, it is all one layer and can be flattened out very easily. There is only one layer of receptors [SHOULD I HAVE A FIGURE OF THIS?]. Regardless of the distance an object is from the eye, the light is imaged by the same receptors, so at this point in the visual system, there is no information about depth. In other words, which receptors actually serve to transduce the light into action potentials does not depend upon how far an object is from the eye. So like motion, extra sources of information about distance have to be used for three-dimensional perception. This information must be found primarily in the visual scene itself, though some will come from sources outside the visual image. In addition, for the first time, the role of having two eyes will become apparent, both in terms of the advantages and the complications having two eyes raise for us.

Depth Cues

Depth perception is an interesting topic for many reasons. On of the interesting reasons is the variety of information in the environment that can contribute to the perception of distance. It is not just one type of feature of the scene but many. In many cases not all of these sources of information are present, but usually there are enough which is what allows depth perception to work in so many situations. These sources of information are called **Depth Cues [to glossary]**. The depth cues can be divided into two main categories: **Monocular Depth Cues [to glossary]** and **Binocular Depth Cues [to glossary]**. The words indicate the difference between these two types of depth cues. Monocular depth cues have the word monocular which comes from the two Greek words {CHECK} mon + ocular which mean one and eye. So, these depth cues only require the use of one functioning eye. Binocular depth cues have the Greek number bin in it which means two. These depth cues require two functioning eyes. Each type of depth cue will be discussed in turn.

Monocular Depth Cues

Go back to Figure 8.1. The photograph is flat. It looks just the same with one eye as with two. This observation should lead to one very obvious conclusion that, however, is often overlooked. The perception of depth does not depend on having two eyes. In fact, many animals rely solely on monocular depth cues, as they do not have any significant overlap in their visual fields from their two eyes (Chapter 4, OTHER CITATION). Their eyes look in different directions, not the same direction like ours. To verify this point about the ability of monocular cues to support depth perception, place a hand over one eye and look around. You can still see that some objects are closer and farther away. More importantly, you can reach out and grab objects that are in front of you. The world has not gone flat. Still, one can find articles about athletes who are supposedly without depth perception and yet they have one perfectly functioning eye. There is an important lesson in this little excursion. Here is an observation, the ability to see depth with one eye that is easy to demonstration and often not noticed. This situation demonstrate both the value of using empirical methods (checking claims by observation) and how little such an approach is used or even understood. So, in this text and elsewhere, where you have the chance, check out the claims.

Back to monocular depth cues. Monocular depth cues have often been broken further into two categories based on whether the cue can be represented in a photograph or painting or not. Those monocular cues that can be represented in a static picture are called **pictorial depth cues [to glossary]**.

The other depth cues cannot be represented in a picture, and for lack of a better term will be called **nonpictorial depth cues [to glossary]**.

Pictorial Depth Cues. For the discussion of the pictorial depth cues, open **Interactive Illustration 8.x, Pictorial Depth Cues [link to media]**. On this illustration you will see 4 circles on a grid, with the top half of the screen a light blue and the bottom half white. To the left are a series of controls that will allow you to add and remove depth cues and manipulate them directly to see their impact on the perception of depth. When the figure opens, the image is flat in appearance. As the depth cues are added, this feature of the image will change.

The first depth cue is **interposition [to glossary]**. In this depth cue, one figure overlaps another. Click on the **Interposition** check box on the upper left corner of the figure. When you select this option, you will see the outer two circles (the blue and red circles) drawn on top of the inner two circles (the magenta and yellow circles). As a result, the blue and red circles appear slightly closer than the magenta and yellow circles. However, the impression of depth is not great. Interposition principally gives information about depth order, or which objects are closer and farther, and not the magnitude of the depth, or how much farther or closer.

The next depth cue derives from the visual angle subtended by an object. If you go back to **Interactive Illustration 4.x, Visual Angle** and reexamine how the visual angle of an object depends upon distance. As a review, the closer an object is, the larger the visual angle of the image on the retina. So generally, closer objects take up more of the back of the eye. The visual system can take advantage of that fact, as seen in the depth cue called **relative size [to glossary]**. If you have closed **Interactive Illustration 8.x, Pictorial Depth Cues** to look at the visual angle figure, reopen it. Also, make sure that the **Interposition** check box is cleared. Drag the **Relative Size** slider back and forth to change the size of the two inner circles. Most people will perceive the two inner circles to appear to move farther away as they become smaller than the outer circles. [I NEED SOME GOOD REAL WORLD EXAMPLES IN PHOTOGRAPHS TO ADDS HERE AND ALL THE OTHER CUES]

Reset the **Relative Size** slider so that all of the figures appear the same size. The next depth cue is called **relative height [to glossary]**. It relates to where an image is relative to the horizon. Be careful, this depth cue is often misdefined even in psychology textbooks. A simple definition is that the closer an object is to the horizon, the farther away it appears. Let us unpack that definition a bit. First take an object on the ground. In this case, the higher up in the scene the person is the farther away it appears. This last statement, shortened to the higher the object is in the image the farther away it appears, is often given as the whole definition for relative height. However, this is only half of the situation. If an object is above the horizon, a cloud, it appears farther away the lower it is in the scene. So the depth cue is relative height. You can demonstrate this fact to yourself on the interactive figure. Use the **Relative Height** slider. As you move the slider up, the circles on the left move down, as if they are on the ground. The outside circle moves lower and appears closer. The circles on the right move up as if they were in the air, above the horizon. The outside moves higher and looks closer. The effect may not be as powerful as for relative size, but most people report seeing it. [DO I WANT TO REFER TO MY ONLINE EXPERIMENT?]. You can see how the depth cue is relative to the horizon, by changing the location of the horizontal line in the figure using the **Horizon** slider on the lower left side of the screen. (You can put the horizon back to its original location with the **Reset Horizon** button at the bottom of the control region on the left side of the window).

Take the next slider, the **Texture** slider, and adjust it to its maximum level, with the slider moved to the top of its range. **Texture gradient [to glossary]** refers to the pattern of the squares on the screen. Up to this time, the squares that make up this pattern have appeared flat. Now the appearance is of walls, a ceiling, and floor going back to a wall that appears farther away. The squares form a texture, and if you follow the squares along the floor they get smaller or finer as they appear to recede. Examine Figure 8.x. This is a photograph of a field of corn. As the corn farther away, the corn gets closer together and it is harder to see the details of the corn. Overall, the texture of the corn is finer. A surface that faces us or is standing up all at one distance, will have the texture be all the same size along its surface, like the back wall in the figure. The texture will affect the apparent distance of the circles. The two inner circles will appear farther away than the outer circles, because they cover portions of the wall at its farthest point. Within the texture gradient image, you can also observe what is often called another depth cue, **linear perspective [to glossary]**. Look at two adjacent vertical lines in the floor of the image. The lines appear parallel, but if you look closely you will see that they converge or get closer together the farther away that they are. You can see the same pattern in the rows of corn in Figure 8.x. However, the vertical lines in the back wall of the image that appear parallel are actually parallel. They do not get closer together and as a result do not appear to get closer to get farther away. [DO I WANT TO HAVE A LINEAR PERSPECTIVE ONLY IMAGE OR A GOOD PHOTOGRAPH].

Look at Figure 8.x, which is a photograph looking across Lake George in Eastern upstate New York. The trees in the foreground look clear and nicely green. The tree covered distant mountains look increasingly bluish. The farthest mountains look little different from the color of the sky. As light travels through the atmosphere, some of it gets scattered. Those wavelengths that we see as blue tend to get scattered the most because of their small wavelength. Thus, the short wavelengths that come from the sun end up entering our eye from all different possible directions, which is why the sky appears blue. Well, what happens for sunlight happens for all light. The farther away the object is the more scattering of its light and the greater chance of blue from other objects entering our eyes. So, distant objects tend to appear a bit fuzzy and bluish. This change in the appearance of object can alter our perception of distance and is called **atmospheric** or aerial **perspective [to glossary]**. You can simulate this depth cue with the **Atmospheric Perspective** slider. Sliding the slider to its **Max Depth** setting blurs the image some and adds in a small amount of blue to the images which does cause a slight perception of a change in distance.

The last monocular depth cue to be discussed is the depth information provided by **shadows [to glossary].** If you adjust the **Shadow** slider, you will see a small gray circle drawn behind the two outside circles (the blue and the red circles). As the slide moves towards the **Max Depth** position, the gray circles are drawn lower and more to the right. What these simulated shadows do is make the red and blue circles appear to move off of the screen towards you. Shadows are an interesting depth cue, in that the use of this depth cue depends upon our interpretation of the direction of the sun. The same shadow pattern can be interpreted visually two different ways, as both an elevation and an indentation. If you look at Figure 8.x, you will see an illustration of the sun shining on both a hole in the ground and on a hill. When the sun strikes a hill or any raised object, the shadow is on the far side of the object. When the sun strikes a hole or indentation, the shadow is on the near side of the object. To interpret a shadow as either being on a hill or a hole, depends upon assuming a direction of the sun and that it comes from above. Look at Figure 8.x (a,b). The first image is the original photograph (a) and it is a picture of an ancient stone carving done in England. These carvings can be found all over England and date back thousands of years. Here is a common pattern of a bowl surrounded by rings. Figure 8.x b is the same photograph but inverted. In this copy of the photograph, all of the indentations look like bowls. The first photograph sets up an expectation for the direction of the sun. Since our expectation is that the sun does not change direction rapidly, the second picture is seen with the interpretation of the same sun direction. The sun is off to the left in the original image, taken when the sun is low in the sky. In the largest of the bowls, the shadow is also to the left, on the near side, as shown in Figure 8.x. When the photograph is flipped over, the shadow is now on the right side, or far from the sun, and it is now consisted with a bump, not a bowl. Thus, our perception is flipped.

With this list, the depth cues available to an artist doing a painting to represent depth have been essentially competed. In doing a painting, an artist will almost never only use one depth cue, but will combine them to create the depth effect that is desired. As more depth cues are added, the impression of depth can be increased. You can simulate how depth cues can be added together, to increase the perception of depth. Along the bottom row of sliders is a slider labeled **All**. By adjusting this slider, you can change the depth setting of all of the depth cues that are checked on the list of checkboxes just to the right of this slider. To save space on the screen, the cues are given abbreviations (**RH** is relative height**, RS** is relative size, **TX** is texture gradient**, AT** is atmospheric perspective, and **SH** is shadow). Select the depth cues you want to add together and use the **All** slider to see if the impression of depth is greater or not than the use of the single depth cues. If you select all of them, you will find that the impression of depth can be rather strong.

[BUILD AN EXPERIMENT TO JUDGE DEPTH ACROSS MULTIPLE DEPTH CUES. SIMPLE VERSION OF PSYCHPLACE STUDY]

Nonpictorial Monocular Depth Cues. There are a couple of depth cues that require only one eye, but cannot be represented in a painting these are **accommodation** and **motion parallax [to glossary]**.

Accommodation is the focusing of the eye by changing the shape of the lens. If you need to remind yourself about accommodation, you can open **Interactive Illustration 8.x: Accommodation** which is a copy of the figure that illustrated accommodation from Chapter 3. With accommodation, different distances from the object focused on will lead to different degrees of effort in focusing. Feedback from this effort might provide some information about the distance of that fixated object. However, accommodation is not terribly accurate. In the daytime, the pupil narrows which reduces the need for accurate accommodation. In the nighttime, our acuity is not very good which also reduced the need for accurate accommodation. So while it is possible that accommodation may be a depth cue, it is probably a poor depth cue.

Motion Parallax, on the other hand, is a very useful and, in many circumstances, very precise depth cue. Open **Interactive Illustration 8.x: Motion Parallax Illustrated**. Motion parallax arises from the motion of the observer in the environment. In this illustration there is one eye drawn, since it is a monocular depth cue, and the eye is positioned on the right side of the screen. The eye can be moved up and down using the **Move Eye** slider on the far right, or by using the **Animate** check box at the bottom right corner of the screen. On the screen are three different dots, objects if you will, each at different distances from the eye. The red object is closest, with the green object at a middle distance and the blue objects is farthest at the left side of the screen. Each dot has a line going from it through the center of the lens to the retina, indicating where the object is imaged on the retina. When the screen first comes up the eye is positioned half way up the screen and in line with the three objects, so that all three objects are all imaged on the fovea. Also when the program first comes up, when you move the eye, it always point in the same direction which is horizontal towards the left side of the screen. Move the eye up and down either using the **Move Eye** slider or the **Animate** checkbox. Notice the motion of the images of the three objects across the retina. The red or nearer object moves to a much greater degree than the green object, which moves more than the blue or most distant object. So, when the eye is focused way off in the distance, the closer an object is to us as we move, the faster it will appear to move past us. Think of driving in a car, and how fast the grass moves by us compared to mountains in the distance. This illustration indicates why the grass appears to fly by.

Use the **Tracking** drop-down menu on the upper left side of the screen. If the menu is in the state it was when the activity started up, it says **No Tracking**. Click on the menu and select **Track Middle**. Now the when the eye moves, it will rotate so that the image of the middle distance object, the green dot, will stay on the fovea. Move the eye back and forth, and observe that the image of the green dot does not move on the retina. However, the images of the other two objects move in opposite directions of each other on the retina. The near red dot moves across the retina in the same direction as the eye moves, while the far blue dot moves in the opposite direction that the eye moves. Recalling that the images on the retina are upside down and reversed anything closer to you than where you are looking, will appear to move in the opposite direction of your movement. Anything farther that where you are looking, will appear to move in the same direction that you are. You can verify that this is the way this motion appears by moving your head back and forth and keeping your eyes still on some object that stands between you and the walls. The walls will appear to move in the same direction as you are, relative to the object you are tracking. Anything closer to where you are looking will appear to move in the opposite direction than the object you are tracking.

By opening **Interactive Illustration 8.x, Motion Parallax**, you will be able see how motion parallax can add to the perception of depth. The figure is similar to the figure used to illustrate the pictorial depth cues, and there are controls on the left side of the screen where you can adjust most of the pictorial depth cues discussed earlier. However, there are a couple of differences. First, there are only three objects, and they are color coded to match the objects in the last illustration. At the bottom of the screen is a **Move Objects** slider and an **Animate** checkbox that can move the objects in the same way as in the last illustration. When the illustration comes up it is in the no track type of situation simulating driving in a car down a road. If you click **Animate** the three circles will move from left to right with the closest, the red circle, moving fastest. You can adjust the other pictorial depth cues and see how the addition of motion adds to the perception of depth. You can also simulate the tracking by using the **Tracking** menu in the lower right corner of the screen. By selecting to track the middle, the green, circle, the green circle will stay still in the middle of the screen. In this case, it simulates you are moving but you are fixating your eyes on this green circle. The motions of the other two dots relative the green dot are exactly the same as without tracking. Remember all motion is relative. The relative motions of three circles are always the same regardless of the way you track your eyes. However, when you are fixated on the middle object, the near, red circle moves in the opposite direction than you move while the distant blue circle moves in the same direction. To best illustrate this point, manually move the **Move Objects** slider. The direction that the slider moves represents your motion. So when the **Move Objects** slider moves to the right, you are moving to the right. When you move the slider to the right, the red circle will move to the left and the blue circle will move to the right. Moving your head back and forth with the **Move Objects** slider might help you get a feel for this situation.

[TO ME: WOULD AN EXPERIMENT BE NICE HERE?]

A depth cue related to motion parallax is **optic flow [to glossary]**. [LOOK UP ABOUT OPTIC FLOW SO I CAN BETTER DECIDE HOW TO TACKLE THIS.] As you move through space, the point that you move toward forms a point of expansion. All of the worlds expands and moves away from that point. Conversely, if you are riding in a car and sitting backwards, the point directly behind you is a point where all the world will disappear into. Open **Interactive Illustration 8x, Optic Flow**. There is a series of squares that grow larger and form a point in the center of the screen. If you click on the **Move** checkbox in the lower right corner of the screen, the boxes will be drawn larger and larger, expanding from the central point. It appears more like a tunnel and creates the perception of moving towards the center. On the lower left corner, there is a **Direction** menu. It is currently selected for **Forward**. If you change the selection to **Backward**, the squares will start shrinking and the motion will appear to be moving away from that point. But still, there is a perception of depth that arises from this motion. [DO I WANT TO TALK ABOUT THE ADAPTATION AT THIS POINT.]

Binocular Depth Cues

To a large extent, the reason our two eyes both look in the same ways is that we can get a lot of information about depth perception from the overlap of the two visual fields. In the area where both eyes see the same part of the world, we have binocular vision. The depth cues that arise from the binocular overlap of the two eyes are particularly powerful. This section will discuss these important depth cues. There are two main depth cues that arise from the binocular overlap of the eyes. Within their range they are powerful, but they also greatly complicate the way the visual system has to process visual space. These complications will have to be discussed and will take some time to explore as these complications have important implications for the visual system. The two depth cues are **vergence [to glossary]** and **binocular disparity [to glossary]**. They will be discussed in term.

Vergence. Vergence is the type of eye movement that was discussed in Chapter 4 where the two eyes move in the opposite direction to each other. Open **Interactive Illustration 8.x: Vergence [link to media]**. This is a simple illustration of vergence. There is a single green object on the screen with lines to each of the eyes on the right side of the screen indicating where the object is imaged on the two retinas. Since this green dot is the object that is being inspected, the image of this green is on both foveas. You can move the object closer and farther from the eye using the **Move Object** slider or the **Animate** checkbox at the bottom of the screen. You can also move the object by clicking on the screen where the object is and dragging your mouse across the screen from left to right. As you move the object, the eyes rotate to keep the object fixed on the two foveas. You can also adjust the separation of the two eyes using the **Eye Separation** slider on the right side of the screen. The eyes have to rotate very accurately to maintain the images on the fovea. Recall that the fovea is the region of our greatest acuity. Small differences of position of the two eyes will be detected, and can lead to double images. So, unlike accommodation, vergence needs to be very precise. Any information about the eye position for vergence can serve as a very precise depth cue for the distance of the object that is being fixated. What is needed is an experimental demonstration that vergence does serve this function.

While there are several studies available to discuss, Leibowitz and his colleagues have done a study that provides a very convincing demonstration. The subjects sat in a dark room staring at a single object. The object was on a black background and positioned at eye level. The darkness and the black background are necessary to eliminate all other depth cues, both monocular and binocular. Recall that in an experiment to demonstrate that your independent variable causes what the experiment finds it is important to eliminate all other possible explanations. The participants had their heads held in place, and they were to fixate on the object. Then the experimenters placed a prism in front of one eye. What the prism does is it moves the image of the object off to one side. It does not change the size (so no relative size) or in any other way change the object. To still see the object as a single object, the eyes must adjust their vergence. If the prism is placed so that the vergence is consistent with looking at an object closer that the actual object, the participants reported seeing the object move towards them. When the prism was reversed, the object appeared to move farther away from the participants (LEIBOWOTZ PRISM REF). Thus, it seems clear that vergence is an effective binocular depth cue. In fact, over the distance that it works effectively, it is a depth cue that has the possibility of registering the absolute depth of objects. However, vergence is only effective for relatively close objects. Go back to the illustration of vergence, and move the eyes as close together as possible. Move the object closer and farther away. If you watcj the motion of the eyes, you can see that there is a much greater change in vergence for objects close to the eyes than farther away. Eventually, objects get so far away, at approximately 20 feet (REF), that vergence no longer changes as the objects get farther away, so it is rendered no longer useful to depth perception.

Stereopsis. The next depth cue to be discussed is binocular disparity, often just called disparity. Disparity is the depth cue, but the ability to see depth via the depth cue of disparity is called **stereopsis** [**too glossary]**. Stereopsis come from two Greek words that mean “solid vision” and refers to the ability of disparity to generate the perception of a three-dimensional world which is characterized by three-dimensional solid objects.

Disparity. Disparity arises because our two eyes are in different locations in our head, and therefore, have a slightly different view of the world. Open **Interactive Illustration 8.x, Disparity [link to media]** to illustrate how the separation of our eyes leads to disparity. In this diagram, you will see the two eyes fixated on a green object with dashed lines indicating the path of the light from this object to both eyes’ foveas. There is a second object in the field of view of these two eyes, the orange object, and the solid orange lines indicate where these objects are imaged on the fovea. On the right eye there is a second darker orange line. This line starts from the center of the lens and intersects the right retina in the same place as image of the orange object falls on the left eye. Indicating where the orange object falls on both the right and left eye indicated in the right eye will help you see the disparity of the images of the orange object. If the darker orange line is on top of the lighter orange line, it indicates that the images for the orange object falls on the same locations of the two retinas. This would be a location with zero disparity. The easiest place to find a location with zero disparity is the fixation point. Either by using the two **Move Disparity Object** sliders or by dragging the orange dot on the screen, place the orange dot on top of the green fixation dot. Now the orange lines are on top of the green lines, and the dark orange line is on top of the light orange line in the right eye. All lines end on the fovea, which is a location of zero disparity if the eyes of the person are properly alinged. Now, use the **Move Disparity Object** slider along the bottom of the screen. Move the disparity object so that it is as close to the two eyes as the interactive illustration allows. First look at the right eye. The image for the disparity object is to the left of the image of the fixation object in the right eye. In the left eye, the image for the disparity object is to the left of the image of the fixation object. If you look at the dark orange line, you will see that it is also to the left of the image of fixation object. The angle formed by the dashed green line and the dark orange line in the right eye is the same as the angle between the dashed green line and the bright orange line in the left eye. This sameness of the angles shows that the dark orange line goes to the same location in the right eye as the image for the disparity object in the left eye. The angle between the dark orange line and the bright orange line in the right eye is the measure of disparity for the two images of the orange object in the two eyes. If you click on the **Animate** checkbox in the lower right corner of the screen, the orange disparity object will move left and right across the screen, towards and away from the two eyes. As it does, watch the relative position of the orange line and the dark orange line in the right eye. As the object gets farther from the fixation object, the disparity of the two images of the disparity object gets larger. You can also drag the disparity object around the screen using your mouse, and see that most locations in the visual world has disparity.

The Horopter. There is a region of visual space where there disparity is zero. This region is called the **horopter [to glossary]**. On the horopter, the two images of an object fall on corresponding points on the two retinas. By corresponding points, it is meant that the brain interprets these locations as having no disparity. Open **Interactive Illustration 8.x: Binocular Disparity and the Horopter [link to media]**. This interactive illustration will look very similar to the last illustration, with two additional features. First, there is a curved blue line that indicates the horopter, and second, there is a **Move**

Fixation Point slider at the bottom of the screen. As in the previous illustration, the disparity object can be moved by either dragging your mouse across the screen, or by using the two **Move Disparity Object** sliders. If you move the disparity object to any location on the blue curve, you will see that the two lines in the right eye indicating the relative location of the two images for the object fall on top of each other, indicating zero disparity. The horopter always goes through the object that is being looked at, so the fixation object is always on the horopter. The horopter shape also changes depending on how far the fixation point is from the front of the two eyes. In other words, the shape of the horopter depends upon the state of vergence. Use the **Move Fixation Point** slider, so that the fixation object is farther from the eyes. As the fixation object moves away from the eyes, the horopter flattens out. Move the fixation object to an intermediate location so that there is both a good distance in front of and behind the horopter where you can move your disparity object. Put it at the same height on the screen as the fixation object for convenience sake, and use the **Move Disparity Object** slider at the bottom of the screen and place the disparity object on top of the fixation object. If you move the object closer to the eyes and then farther from the eyes than the horopter, you will see that the disparities around the horopter form a mirror image of each other. The disparity is small near the horopter and gets larger as the object moves away in either direction. The only difference between disparities closer or farther than the horopter is the relative position of the two images. Looking at the right eye, when the disparity object is closer than the horopter, the right eye image is to the right of the image of the fixation object. When the disparity object is farther than the horopter, the right eye image is to the left of the image of the fixation object. Thus, the relationship between disparity and distances is tied to our state of vergence, which determines the location of the horopter.

Since the magnitude or size of disparity does not tell whether the object is in front of or behind the horopter, an additional distinction is needed. A simple demonstration using your hands will most clearly illustrate this concept. Take one hand, say your right, and with your forefinger up put it very close to the front of your nose, but where you can still make out your finger. Take your left hand, and extend it to arm's length but straight in front of your other hand. Look at the close hand's finger. The horopter is going through that finger, and the finger of the other hand lies way behind the horopter. The left hand's finger is so much behind the horopter that you should see two images of that finger. Stay looking at the close finger and close your right eye. The distant left finger should appear to the left of the closer right finger. In the left eye, the distant left finger is to the left of the object being fixated. Next, close the left eye and open the right eye. Now the left hand finger is on the right of the right hand finger. Since the disparate object, in this case the left finger, is on the same side relative to the fixation object as the eye in which the images are falling, this situation is called **uncrossed disparity [to glossary]**. The opposite is true if you fixate on the distant finger. Now close the right eye. The closer right hand finger, which is the finger seen in disparity, is now to the right of the fixated finger, which is the left hand finger. The relative positions of the fixated and disparate images are reversed. If you open the right eye and close the left eye, the right hand finger will now appear to the left of the fixated finger. These are crossed disparities. Recalling this demonstration, open **Interactive Illustration 8.x, Types of Binocular Disparity [link to media]** to further clarify why disparity works this way. This interactive illustration will show you the relative position of the two images on the two eyes. This illustration is similar to the last illustration, but again there are two new additions. First, there is the same fixation object, horopter, and disparity object that was there in the last illustration. You can move the disparity object closer and farther from the eyes and change the eye separation. On the left side of the screen are two new circles that illustrate the relative positions of the images of both the disparity and fixation objects on the retina. The top circle, labeled **Layout on the Retina** shows how the images from the disparity and fixation objects layout on the two retinas if they were overlapped centered on the foveas. So, the fixation object, shown as an **F** on the screen, shows only one image as both images are on fovea and the two foveas are in the same location. The disparate object, in most cases, will have two images so each retina's image is indicated by the letter for the corresponding eye, **L** for the left eye and **R** for the right eye. Now the images on a retina are upside down and backwards, so the letters are written upside down and backwards. The bottom circle, labeled **As We See the Images**, shows us how we experience these images since we do see the world upside down and backwards. This circle is simply the top image rotated so that the world appears upright and forwards. When the illustration first comes up, the disparity image has an uncrossed disparity, that is, it is behind the horopter. If you look in the **Layout on the Retina** figure and compare it to the layout of the images in the two eyes, you will see that the image coming from the right eye is to the left of the image of the fixation object in both cases. In both cases as well, the image of the disparity object is to the right of the fixation object in the left eye. It is particularly easy to see the resemblance between the **Layout on the Retina** circle and the image by just looking at the right eye where the position of image from the left eye is superimposed. When you look at the **As We See the Images** figure, you will see that the left eye image, **L**, is to the left of the fixation point, **F**. That is, the disparity object image is to the same side of the image of the fixation object as eye that the images are in, and, thus, these images are said to be uncrossed. Use the **Move Disparity Object** slider, or drag the disparity object to in front of the horopter. Now the situation regarding the relative positions of the images of the fixation and disparity objects has reversed. On the eyes and on the **Layout on the Retina** figure, the image from the disparity object is to the left of the image of the fixation object. However, in the **As We See the Images** figure, the left eye image for the disparity object is to the right of the image of the fixation object. As a result, this is called crossed disparity.

Does Disparity Work? Now the question remains: does disparity actually generate the perception of depth when it is the only depth cue? Open **Interactive Illustration 8.x, Stereopsis [link to media]**. To use this illustration it is necessary to use the glasses with the colored eye pieces. Place them on so that the red color is over the right eye (to help you remember, both red and right begin with r). When you put the glasses on, the colored eye pieces act as filters. There are actually two circles drawn on the screen, a red circle and a blue circle. The red eye piece pretty much limits all but red wavelengths of light into your right eye. The blue eye piece limits all but blue light into your left eye. So each eye gets mostly only one of the two circles, with the red circle going to the right eye and the blue circle going to the left eye. Initially, the two circles are drawn in the same place on the screen. Take off the glasses for a minute and move the **Disparity** slider to one end or the other and you will clearly see the two circles. Move the **Disparity** slider back to the middle and replace your glasses. This time, slowly move the disparity slider to the right and the top circle will appear to move forwards and the back circle will appear to move backwards. Move the slider only a couple of steps. You can lift or remove your glasses and see that only a very little disparity is needed to be able to see the change in depth. Stereopsis, like vernier acuity, is a hyperacuity. A disparity of only about 7-8 arcseconds of a visual angle is all that is needed for stereopsis to be perceived (REF, XXXX). Recall that the diameter of cones in the fovea, where they are the smallest, is about 30 arc seconds. So a disparity that is less than the diameter of a cone is all that is needed to cause the perception of depth. You can select different types of shapes to examine, by using the **Shape** menu down in the lower left corner of the screen.

[WHAT ABOUT DOING SOME PICS – PERHAPS SOME FROM THE MARS ROVER IF I CAN GET PERMISSION]

Now open **Interactive Illustration 8.x, Random Dot Stereograms [link to media]**. Before you put on the glasses you will see two overlapping patterns of dots, one red and one blue. When you put on the glasses, most of you will see a smaller square in the middle of the field of random dots jump out and appear to be floating in front of them. Some people take a while to see this pattern, and a few cannot see the square at all, so if you cannot see the square, give it a little while but do not be concerned if it does not pop out (Ref, XXX). The red and blue patterns of dots are nearly identical. The dots that form the background are the identical, but offset from each other by the small degree indicated by the part of red and blue dots that do not overlap at the left and right edges of the figure. The dots that form the square are also identical, but these dots are offset from each other in the opposite direction as the background dots, making the disparity crossed so that the dots float in front of the background. There are several remarkable features of this illustration, which was first explored by Bela Julesz (REF, XXXX). First, just as common motion can generate the perception of a rigid form, so can differences in depth. The floating square appears to be a coherent shape, but the only defining feature is that it has a different disparity than the background. There are no contours or even motion differences to form these edges. Second, the rapidity with which the square is seen against the background is quite remarkable. The dots that make up the stimulus are really indistinguishable from one another, and yet the brain can quickly match the appropriate dots from the two eyes to determine the disparity of each eye. To show how remarkable the speed with which the dots can be matched to determine disparity really is, try a demonstration to push this matching ability. Click on the

Dynamic checkbox in the lower right corner of the screen. Then, use the **Dot Size** slider at the bottom of the screen to increase the size of the dots. Increase the dots size so that the pattern of dots for the red and blue images update at a fairly fast pace. Regardless of the speed of the update, if you can see the square, usually you will still see the square even though the pattern of dots making up the square changes continuously and randomly. The brain has no trouble find the proper mate for all of these identical dots (REF – K White?, XXXX).

The Structure of Visual Space. Reopen **Interactive Illustration 8.x, Stereopsis**. Put the glasses back on and start with the disparity at zero. Then gradually move the **Disparity** slider to the right, so that the top circle moves towards you. As you slowly increase the disparity, the top circle moves forward off of the screen. However, as the disparity continues to increase, you will begin to see the edges of the two circles through your glasses and eventually the circles may even separate and no longer blend into a single circle. There comes a point when disparity becomes too great for the visual system to handle and the images on the retina are no longer blended together into a single perception. Open **Interactive Illustration 8.x, Panum's Fusional Area**. This figure is similar to **Interactive Illustration 8.x, Binocular Disparity and the Horopter**. There is one additional feature. Around the cyan horopter line, there is a brown field drawn. This field represents **Panum's fusional area [to glossary].** Panum's fusional area is a region that is centered on the horopter where the disparity sufficiently small enough to allow for the two images to be fused into the perception of a single image. Use the **Move Disparity Object** slider to move the disparity object to the fixation point (in this illustration you can only move the disparity object closer and farther away from the eyes for simplicity's sake). Then move the disparity object towards one end or the other of Panum's fusional area. When you reach one edge, look at the amount of disparity for the images in the eyes. This disparity is the greatest amount of disparity that the eyes can fuse into the perception of a single image. If you move the disparity object to outside of Panum's fusional area, the brain cannot fuse the two images and double images or **diplopia [to glossary]** results. The absolute physical size of the region depends upon the observer's state of vergence just like the horopter does. When the eyes are fixated on an object close to the face, the rate that disparity increases as an object moves away from the horopter is very fast. When the eyes are fixated on a distant object, disparity increases at a much slower rate when as an object moves from the horopter. You can illustrate this feature of spatial vision by using the **Fixation Point** slider to change the distance of the horopter from the eyes. After you have a new location for the fixation point, you can use the **Move Disparity Object** slider to see how the disparity changes as the object moves away from the horopter. Actually, the visual system can only handle a small range of disparity. The limit of Panum's fusional area is about 32-40 arcmin (Qin, Takamatsu & Nakashima, 2006). The size of Panum's fusional area is 5 to 7 times the actual size on the screen. [GET THIS VALUE].

Diplopia and Binocular Suppression. Given the small size of Panum's fusional area, most of our visual world is in a region where the two images cannot be fused. The two images stay separate, and we experience double images or **diplopia**. You experienced diplopia in the demonstration where you held your fingers in front of your eyes to demonstrate disparity. When you fixated on the finger in front of your face, you could see two copies of the distant finger. The distant finger was outside of Panum's fusional area for fixating on the close finger. This explanation may raise a very important question. If Panum's fusional area is small and most objects in our visual world generate double images that cannot be fused, why are we not more aware of double images?

Try a simple demonstration if you have the proper materials handy. Get a paper towel role, or any other long tube. Put the paper towel role up to one eye and look through the tube, keeping both eyes open. Put the hand corresponding to the eye that does not have the tube at the end of the tube. Place it with the fingers up and the palm facing your face. For example, if you have the paper towel tube over your left eye, put your right hand at the end of the paper towel tube in front of the right eye (Figure 8.x). Now, look at something in the distance, say a wall at the far side of a room. The fovea of each eye will have a different image. If you lay it out as above, the left eye will have the wall falling on its fovea and the right eye will have the palm of your hand on its fovea. This is a form of diplopia. However, you will not see a confusion of the two images or a blending, but you will see the wall and what appears to be a hole in your right hand. Since you are looking at the distance, focusing on that, the image of your hand is suppressed and you do not experience it. This is known as **binocular suppression [to glossary]**. The image of one eye is not experienced.

This demonstration is a very powerful demonstration of binocular suppression, and it is strongly recommend that you try it. However, it is quite possible that you do not have a paper towel tube handy and there is more to be observed here than just binocular suppression. So open **Interactive Illustration 8.x, Binocular Suppression and Rivalry [link to media]**. When you open the image, you will see a relatively large blue square covered by a textured red circle where the texture is changing periodically. If you put on the glasses, you will see that the blue square goes to the left eye and the red circle to the right eye. When you close your right eye, you see a smooth blue square with out the hole caused by the red circle. When you close your left eye, you will only see the red circle. When both eyes are open, you will see a hole in the middle of the blue square caused by the red circle. Here is another illustration of binocular suppression.

Suppression relates to stereopsis in dealing with the diplopia for objects outside of Panum's fusional area. While most of the visual world is made up of objects that generate double images, in most cases, one of the images is suppressed. That way we only are aware of one of the images and are not disturbed by the double images. There are limits to suppression as was experienced when looking at a finger right in front of your face and seeing both images of the finger held at arms length.

It is even the case that sometimes the brain cannot decide whether to use the right or left eye image. Use the **Binoc Stim** menu at the lower left hand corner of the screen and select the **Rivalry** option. Keep your glasses on or put them back on to your face. This stimulus is two distinct gratings. There is a red vertical grating going to your right eye and a blue horizontal grating going to your left eye. You can verify the stimulus by alternately closing each eye. If you remove your glasses, you will probably see a checkerboard. With your glasses on, the situation is quite different. The two stimuli are not blended, but they compete with each other. You might see the blue grating or the red grating. You might even see in some places one of the gratings and in other places the other grating. On top of this, the experience is not static. You might have started seeing the red grating, but it will change to the blue grating and back. But you will not get much of a sense of a checkerboard. The brain does not add the two stimuli together but tries to pick one to see, but there is not obvious advantage for one stimulus than the other so the situation is dynamic and constantly changing. This is **binocular rivalry [to glossary]** and is what happens when binocular suppression or fusion is not possible. As you have experiences, binocular rivalry is not pleasant and is a good experience to avoid.

To summarize, having two eyes that look the same direction gives our visual system some powerful new depth cues: vergence and disparity. However, two forward looking eyes also complicates the nature of the visual information that the visual system needs to handle. There is a limit to the range of disparity that our visual system can handle completely. This limit leads to Panum's fusional area which defines the region of visual space where complete fusion of the two disparate images is possible. Panum's fusional area is quite small. Outside Panum's fusional area, fusion is not possible. We might experience double vision or diplopia. However, in most cases, one of the two images is not seen but is suppressed. If suppression is not possible, binocular rivalry might sometimes result where there is an active competition between the images of the two eyes to determine what is seen. The limited range of fusion leads to the need to actively handle the diplopia that results. The precise depth cues are bought with a price in complexity.

I. The Accuracy of Depth Perception
II. The physiology of Depth Perception

Figure 8.1. The Rocky Coast of Maine. See all of the objects at different distances from you on this flat picture.

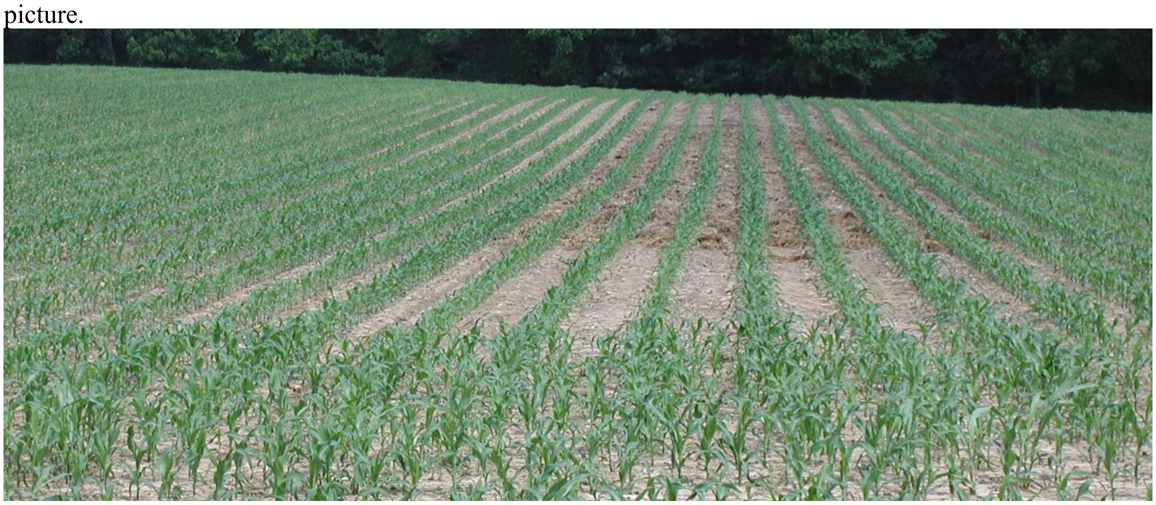

Figure 8.x, A cornfield. Notice how the corn gets closer together in the distance.

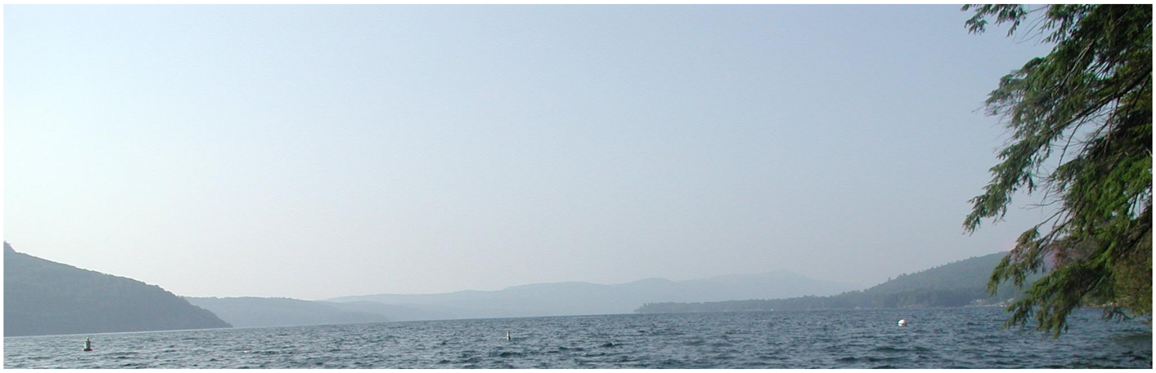

Figure 8.x. Photograph showing atmospheric perspective.

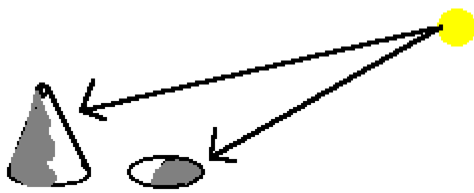

Figure 8.x An illustration of how shadows relate to the direction of the sun.

(a)

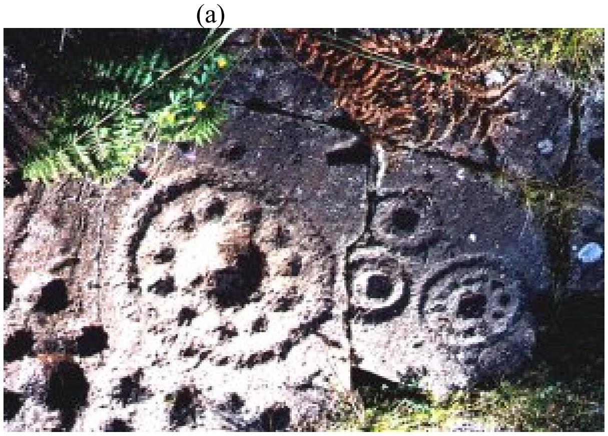

(b)

Figure 8.x (a) a photograph of an ancient stone carving in England. (b) same picture as in a just flipped upside down. Notice how all the indentations now appear as bumps. [I STILL NEED PERMISSION FOR THESE PHOTOGRAPHS].

Figure 8.x. Illustration of how to do the paper towel tube demonstration of binocular suppression.