Chapter 5

Contour and Form

Chapter Outline:

- I. Introduction
- II. Seeing Differences
 - a. Lateral Inhibition and Receptive Fields
 - b. Center-Surround Receptive Fields and Lateral Inhibition
 - c. Role of Contour in Perception
 - d. Contour and Contrast in Application
 - e. Limits to Lateral Inhibition as an Explanation
 - f. Feature Extraction
- III. Patterns Across Regions
 - a. Spatial Frequency and Sine Wave Gratings
 - b. The Contrast Sensitivity Function
 - c. Fourier Analysis
 - d. Bring the Contrast Sensitivity Function and Fourier Analysis Together
 - i. The Concept of Channels
 - ii. The Blakemore-Sutton Effect
 - iii. Adaptation and Aftereffects as a Research Tool
 - iv. Spatial Frequency Channels and Adaptation
 - v. A Blakemore-Sutton Effect Experiment
 - e. Physiological Evidence for Spatial Frequency Channels
 - f. Applications of the Contrast Sensitivity Fuction
- IV. Grouping and Gestalt Contributions to Form Perception
 - a. Background
 - b. Figure-Ground Perception
 - c. Grouping and Organization
 - i. Laws of Proximity and Similarity
 - ii. Law of Good Continuation
 - iii. Law of Closure
 - iv. Law of Symmetry
 - v. Law of Common Fate
 - d. Texture Segmentation
 - e. Good Figure or Prägnanz
- V. Higher Level Processes
 - a. Top-Down vs. Bottom-Up Perception
 - i. Ambiguous Figure-Ground Percpetion
 - ii. Expectations or Perception Set
 - iii. Context, etc?

Introduction

What has been covered so far about the visual system can be summarized as the basic anatomy, physiology and abilities of our visual system. For instance, the last chapter covered some of the fundamental abilities of the visual system. In particular, a fundamental ability of the visual system is its acuity. But the ability to measure our acuity assumes numerous visual capabilities that were simply glossed over in the last chapter. Take the Landolt C image that was illustrated in the last chapter (Open Media Figure 5.x, Landolt C, which is a repeat of the media Acuity Stimuli Figure from the last chapter). The feature of this image that was relevant for acuity was the ability to locate the gap in one of the sides of the circle. But locating the gap requires perceiving the **C**. Before you can do that, you need to see the curve making up the **C** as different from the background. This is a process sometimes called segregation (Beck, 1993; Spillman, 1999). Now notice the checkerboard. The squares, while seen as separate squares, are also seen as all belonging together. It is the perception of the unity of the squares that allow us to apply one word to the stimulus, "checkerboard". Parts of an object that belong together need to be perceptually grouped and seen as a unit, even thought they stimulate separate receptors and retinal ganglion cells. This perceiving separate parts of an object as a single stimulus is called grouping (Kubovy & Wagemans, 1995; van Lier & Wagemans, 1997, 1998). Beyond these basic abilities, it is usually helpful to identify the object (Mishkin, Ungerleider, & Macko, 1984; Sanocki, 2001). These examples highlight some of issues involved in form perception, which is the topic of this chapter. This chapter is divided into four parts. The first part deals with the most basic aspect of form perception, the ability to detect differences in the world around us. The second part of the chapter will cover the contrast sensitivity function and what it tells us about how we perceive the world. The third section deals with what could be called intermediate levels of form perception, which involve perceptual functions like the grouping mentioned above (van Lier & Wagemans, 1997, 1998). The final section is related to what has been called higher functions in form perception, which deal with object identification (Mishkin, Ungerleider, & Macko, 1984; Sanocki, T, 2001). Many of the most important features of form perception cannot be predicted from what happens at the retina and first stages of processing in the cortex. It is these interesting phenomena that will be discussed in the last section of this chapter.

Before we begin, let us note that a couple of simplifications have been taken in this discussion of form perception. First, the forms will all be two-dimensional. Real objects are usually three-dimensional. Second, all of the issues discussed in this chapter will related to stimuli that change in luminance, and so all of the examples will be in shades of gray. Both of these topics will be covered in later chapters. You can think of the organization of these chapters as adding complexity. We talked about the most basic functions last chapter. Now we add the discussion of how we see objects that change in luminance in only two-dimensions. Next chapter, color will be added to the visual world and the chapter after that, the bounds of two-dimensions will be broken.

Seeing Differences

The first place to start a discussion of form perception is with our ability to detect something, anything, against a background. How do we know that there has been a change and what does it take for that change to be registered? The discussion of acuity in the last chapter provided part of the answer – the change has to be big enough to be registered. But that is not all. If you still have the acuity stimuli figure open, notice that every place that you see something, the brightness changes. Even the luminance changes on the screen. As you will see later, there is a very great difference between luminance and brightness. These changes in brightness, or contours, are very important for our form perception.

A **contour [to glossary]** is an edge. More precisely, it is where there is a fairly abrupt change in the visual stimulus that allows us to see an edge. Contours define the limits or outline of an object, such as the outline of a face. In addition, these contours may define the outline of some feature of an object, such as the eyes of a face. These contours are important guides for our eye movements. When we view a scene, we tend to move our eyes to regions of the scene that are rich in contours (Buswell, 1935; Mertens, Siegmund, & Gruesser, 1993). The changes in the visual stimulation that can create contours can be based on changes in luminance, color, texture, motion, and even depth. In this chapter, we will discuss contours that arise from luminance and texture. First, luminance contours will be discussed, and then under grouping the texture contours will be discussed.

So for now, a contour is a change in the luminance or light level from two adjacent regions of the visual world. Open Interactive Illustration 5.x, Basics of Contours [link to media]. This figure appears as a grey field when you first bring up the screen. Actually, there is a square in the center of the screen that you cannot see, because there is no contrast or contour with the surrounding grey. Use the slider on the side to adjust the luminance of the central square. If you move the slider up, the luminance increases and the square begins to lighten and eventually becomes white. If you move the slider down, the square appears darker and eventually black. Adjust the luminance of the square to one extreme and then reduce the contrast between the square and the surround by adjusting the slider so that the difference in the luminance between the square and surround gets smaller. As you reduce the contrast, you will reach a point where the contour between the surround and the square are below your threshold. At that point, the square becomes invisible. So now we have a new type of threshold, the threshold of luminance differences that is necessary for a contour to be perceived. You can also change the shape of the figure being drawn. You can have a circle or an edge as well, by using the drop-down menu in the upper left hand part of the screen where you see the word **Square**. Notice that in these very simple figures, the shape is determined by the shapes of the contours. If you choose either the **Edge** or **Bar** shape, you can make the contour either **Sharp** or **Sloped**. The sharp edge is what you have with the square and circle. The **Sloped** edge is a blurred edge. See if it takes the same amount or a different amount of contrast to be able to see these figures with the sloped or sharp edges.

Given the importance of luminance contours, it is not surprising that there are ways to measure contours. Since for the current purposes only contours created by changes in luminance are being discussed, these contours are said to occur as the result of **contrast [to glossary]**. To be able to see a contour created by luminance, there needs to be sufficient contrast. The simplest way to measure contrast is by the contrast ratio, or simply the ratio of the luminance of the brighter of the two regions (Lum_{bright}) divided by the luminance of the darker region (Lum_{dim}).

$$Contrast = \frac{Lum_{bright}}{Lum_{dim}}$$

There are many other measures of contrast that are useful for in research (e.g., Morse, & Rosenthal, 1996; Switkes, & Crognale, 1999), but the contrast ratio is the easiest to comprehend, and so will be used here. An example from the application of vision in the Human Factors area will illustrate this point nicely. One reason this example is a good one, is that in the applied world the contrast ratio is a fairly common measure of contrast. For example, while we can make out objects with a very small contrast ratio, we will be faster at making out the objects as the contrast ratio increases, up to a contrast ratio of about 3 to 1 (Krantz, Silverstein, & Yeh, 1992; Other ref). That is where the brighter region is three times more luminous than the dimmer region. However, having a contrast ratio of greater than 3 to 1 will not make responses faster or more accurate. Yet, when people, in this case pilots, are given freedom to set the contrast ratio of a display, they prefer about twice as much contrast ratio (Silverstein & Merrifield, 1985). To help place these values in context, a standard television has a contrast ratio of about X to X (REF).

Sufficient contrast in an image to see these contours is very important to our ability to perceive an object. Open up **Interactive Illustration 5.x:** Contrast and Form [link to media], to see a simple illustration of the important of contrast to the ability to discern the contours in an image and, thus, to be able to see the objects. When you open the image, you will see a picture that does not reveal a lot. Imagine an image yourself on a foggy day. Fog reduces differences, contours, and obscures much of what you can make out. In this figure, the contrast (that is, the differences between the pixels) has been reduced to a very low level. You might make out some aspects of the photograph. Slowly drag the **Contrast** slider down and the contrast will gradually increase, so that you can make out the image. What allows you to make out more and more of the scene is that as the contrast increase, your visual system is able to detect more of the contours contained in the image. At first, you might make out some of the nearer objects, such as some of the trees and bushes. As the contrast increases, gradually you will see the Ohio River and hills in the background. The details of the sky and cloud, which have the lowest contrast, will be the last to emerge.

So we need contrast to be able to detect contours and, thus, form. Let us do a demonstration of a couple of classic visual phenomena that illustrate this fact. Open <u>Interactive Illustration 5.x: Contour</u> and Form [link to media] for us to illustrate how contours can drive our perception. When you open the figure you will be faced with a personal favorite. This illusion is called the Craik-O'Brien-Cornsweet illusion (Cornsweet, 1970). You think you are looking at a circle in the middle of the screen. Yet, the

[1]

center of the circle has exactly the same luminance as the outside edges of the gray display area. The only changes in the luminance of the screen are at the edges. There is a dark contour on the outside of the circle and a light contour on the inside, and these contours create your perception of the circle. I challenge you to prove this claim to yourself. Cover the edge of the circle and see if the effect does not go away. The entire illusion is created by the edges. You can control the intensity of the edges with the **Intensity** slider and the width of the edges with the **Edge Size** slider, both on the right hand side of the screen. As you decrease the intensity, the effect will not be present for the lowest intensities. As you increase the intensity, you will find a place where the effect is optimal, in that the circle looks the same all the way across. With further increases in intensity, the illusion will remain but the edges will become visible. As you reduce the size of the edges the effect reduces and eventually disappears. See if you can find the optimal effect on your monitor. If you press the **Direction** button, you can change the direction of the edges and make the illusory circle now darker than the background. Pressing the **Plot Contour** check box will show a line graph across the screen of the intensity of the figure. At each point, the relative height of the line will indicate the relative intensity of that point on the screen. At the edge and the center, you will see the plot will be at the middle, showing that they are the same. I still think the best proof is to cover the edges. If you keep the plot on your screen and change the image, the plot will keep up with your changes.

So, the Criak-O'Brien-Cornsweet illusion is a demonstration of how contours can create the perception of an object that is not there. First, press the **Direction** button so that you have a bright center circle in the Craik-O'Brien-Cornsweet illusion. Where it says **Craik-** in the upper left hand corner of the screen, there is a drop menu and you can select another illusion that is derived from Cornsweet (1970) but does not have a fancy name. It is called **Minimal Contours** on the menu. Select it. You may see a small bright region, in the screen but actually the circle is quite a bit brighter than you suspect. First, reduce the edge size and watch the central circle appear. The intensity does not change at all for the central region. You can verify this fact yourself by using the **Plot Contour** feature again. Here by not having good edges, blurring them, an object can be rendered invisible. It is not enough to just have contrast. In essence, the contrast, as measured by the ratio of the luminance of the bright region over the dark region, has not changed because these luminances have not changed. You also need a sharpness to the change in luminance, that is, you need a contour. You can use the **Change Color** button below the **Plot Contour** button you can have the image drawn in different colors to see if color affects the illusions. At this point, it is time to look again as some of the features of the retina and receptive fields to see if this role of contours in perception can be explained.

In both of these examples, the perception of the central region is determined not by its own luminance but by the luminance surround it. In the Craik-O'Brien-Cornsweet illusion, the luminance at the edge of the circle is filled in to the center. In the minimal contours illusion, the luminance of the surround fills in to the center because there is no contour to stop it. Open Interactive Illustration 5.x, Filling-in to see another illustration of this phenomena of filling in that occurred these two demonstrations. When you open the figure, you will see a blurred image with a black fixation point in the middle. Stare at this fixation point for about one minute and observe what happens to your perception of the blurred circle. It is very important to keep your eyes as still as possible during this exercise. Try it out and then return to the text.

For most people the circle will nearly disappear, and for some it might disappear completely. This example is similar in a way to what was experienced in the minimal contours, only much more slowly. In this figure, filling in again happened because of a lack of contours as in the minimal contour illusion. You can manipulate the level for the drawing of the center of the blurry circle, the outside of the circle and the size of the circle to see how this affects the ability of the filling in mechanisms to make the blurry circle disappear. Filling in has seemed to occur in a few situations now: with the blind spot, minimal contours, and now this last illustration.

Lateral Inhibition and Receptive Fields

To better understand the special role of contours in our perception, it is necessary to reexamine the concept of receptive fields, in particular the receptive fields of the ganglion cells. These center-surround receptive field shapes do something very interesting to contours, especially contours that are edges like in **Interactive Illustration 5.x, Basics of Contours**. But before we get to these receptive fields, let us simplify the situation and look at a much simpler visual system. To explain what these receptive fields may

be doing, let us look at vision in another animal, the **limulus [to glossary]** or horseshoe crab (Figure 5.x). Ratliff and Hartline (1959, Ratliff, 1971) used this animal to do some fundamental research on how the eye processes visual information, that still influences how we think about vision today. This work led to Hartline receiving the Nobel Prize in 1967. The limulus has a compound eye much like the fly, as shown in Figure 5.x. The compound eye is made up of many different little eyelets called ommatidia. Each ommatidia is complete with its own cornea and receptor. Given the size of the ommatidia, it was easy to put light into one of them without having any of the light falling on any of the adjacent ones. This feature is crucial to the research about to be discussed.

Open up **Experiment 5.x, Lateral Inhibition [link to media NOT DONE JUST POWEROINT VERSION FOR NOW TEXT FOLLOWS FUTURE FIGURE]** and we will perform a simplified version of one of Hartline's experiments. In the center of the screen, you see a set of idealized ommatidia. They are labeled **A** through **X** for discussion purposes. Above each of the ommatidia there are two small buttons: **S** and **C** for stimulate and clear. Pressing the **S** button will cause a beam of light to proceed to the ommatidia. The strength of the stimulation will be indicated by a number above the **S** button. Repeated clicking on the button will alter the strength of the stimulation. Below each of the ommatidia is a number indicating the firing rate of that one. In the actual experiment, these recordings were done by having a microelectrode penetrate the nerve leaving each ommatidia.

When you first bring up the simulation, it will be in the dark. Like any good neural cells, they have a low firing rate in the dark. However, to simplify this outcome, there is no variation in their firing rate like there really would be. What your task will be is to try to understand what happens to the firing rate of one ommatidium when you stimulate the next door ommatidia. So pick an ommatidium, say **C**, and press the **S** button for that ommatidium. It really does not matter which one. Notice the increase in the firing rate of the cell. Next, add some light to one of the ommatidia that are adjacent to the first ommatidium you stimulated. Press the S button for that cell. The trick is not to examine the firing rate for the new ommatidium, but the firing rate for the ommatidium you first stimulated. Say the first cell is **C** and the second cell is **D**. Just examine the firing rate of **C**, while you add and remove the light stimulating **D**. Notice that as you add light to **D** and do not change the light to **C**, you reduce the firing rate on **C**. Adding light to an adjacent cell is called lateral inhibition [to glossary]. Try increasing the stimulation to ommatidium **D**. Notice that as you add more light to **D**, the firing rate on **C** continues to decrease. The amount of lateral inhibition that one ommatidium generates is proportional to how strongly it is stimulated. Lateral inhibition may be very counterintuitive. Why should light in one ommatidium affect the response of the neighboring ommatidium? It turns out that this lateral inhibition is tied up very intimately with the role that contours have been playing in our discussion.

Let us use this simulation to do a simplified version of a very important experiment performed by Ratliff and Hartline (1959). First, clear all of the ommatidia that you have stimulated so that we start at a common point. Then, stimulate the left half of the ommatidia by pressing their **S** buttons 4 times so that they have a stimulation level of **20**. Now for the cells on the right half of the screen, press their **S** buttons 2 times so that they have a stimulation level of **10**. The crucial cells to examine are the two cells at the edge you have formed. You can think of this stimulus as a light edge on the left, with a dark edge on the right, much like the **Edge** stimulus that you could build in **Interactive Illustration 5.x**, **Basics of Contours**. So find the cells at the edge between your two fields. Now, the light level on the left is constant (make sure it is) on the left half of the field. However right at the edge, the last cell on the left actually has a higher firing rate that the rest of the cells on the left. Conversely, the cells on the right all have the same stimulation, but the cell at the edge actually has a lower firing rate than the rest of the cells on the right.

The edges of a stimulus, where the contour is, is actually enhanced by the lateral inhibition. The ommatidia that receive the white light generate comparatively a lot more lateral inhibition that the cells that receive the light from the gray patch. So an ommatidium in the middle of the white patch will receive a fair amount of inhibition from the ommatidia near by.

[how quantitative should I make this illustration]

However, an ommatidium stimulated by the white patch, but near the contour between the white and gray, receives much less inhibition because part of its lateral inhibition comes from the cells that are stimulated by the gray patch. Thus these cells near the contour respond more strongly than cells in the middle of the white patch.

Conversely, the cells receiving light from the gray patch generate less lateral inhibition. But the cells stimulated by the gray patch that are near the contour with the white patch, receive more overall inhibition because of the inhibition generated by the cells stimulated by the white patch. Recall that the amount of lateral inhibition generated by an ommatidium is proportional to the strength of the stimulation.

While it is interesting that the limulus has lateral inhibition, you might be wondering what all this has to do with the way that humans see form. Let us start with a subtle but well-know visual phenomenon that seems to suggest that lateral inhibition operates in our own visual system. Open Experiment 5.x, Mach Bands [link to media]. This is an interactive illustration of the famous Mach Bands, first described by the physicist Ernst Mach, who also gave us one of the first descriptions of color blindness, to be discussed in the next chapter. What Mach found is that our perceptions of the brightness of objects do not always match the changes in luminance. Remember that brightness refers to your subjective perception of intensity, and luminance is the physical measure of intensity. This will be a very simple magnitude estimation of brightness. On the screen is a blurred bar, much like the blurred bar in the **Interactive Illustration 5.x, Basics of Contours**. On the left is a slider that you can use to adjust the intensity of the bar, but do not do anything with it right now. You are first going to do a quick magnitude estimation experiment, using the **Adjust Dot Brightness** slider between the bar and the open graph in the right half of the screen. Below this slider is a button that says **Match**. There is also an arrow on the screen pointing up. At that location, move the Adjust Dot Brightness slider up or down until the brightness of the dot below the arrow matches your subjective impression of the brightness of that location on the screen. The highest value of the slider will make the dot completely white, bright white, and the lowest value of the slider's range makes the dot the blackest possible. So as you adjust the slider, know what your limits for adjusting the slider are and give yourself room to adjust to brighter and darker perceptions. When you are done with this judgment, press the **Match** button and the arrow will go to the next location, and you will repeat the procedure. You will do this for 15 locations across the screen. When you are done, the Match button will be deactivated. Complete the task and then return to the text.

Click on the **Luminance** check box below the graph. When selected, this check box will show both your results and the relative intensity of the stimulus across the screen, scaled to match your judgments. Notice that the two are quite similar but they show some odd behavior at the edges. At the edge of the bright region you perceive that the area is brighter than it should be for the edges and at the edge of the dark region you will find that you judge the edges darker than they ought to be. Given the fairly universal nature of this phenomenon, I feel confident in making these predictions. The extra bright and extra dark regions follow the edges, so they form thin bars or bands.

These illusory bars are the Mach Bands that give this illusion its name. If you refer back to the results of the light and dark region in our **Experiment 5.x: Lateral Inhibition**, this is essentially the result that we found with lateral inhibition. When the light and dark regions were placed adjacent to each other, the output of the ommatidia showed a pattern of response that was similar in pattern to what you found in our Mach Band experiment. This suggests that there might be something similar to lateral inhibition going on in the eye. The question is now where does this lateral inhibition occur in the mammalian visual system? The center-surround receptive fields seem to be a very good candidate for this type of response.

Center-Surround Receptive Fields and Lateral Inhibition

If we reexamine the center-surround receptive fields, we will see that they are very well suited for producing the kinds of effects that were seen in the Mach Bands. Recall that the center-surround receptive fields have both excitatory and inhibitory regions to them. Open **Experiment 5.x Lateral Inhibition in Receptive Fields [link to media]**, and let us take a closer look at these receptive fields. Our goal will be to deepen our understanding of how these receptive fields impact the way we see. In this experiment, we will be working with an on-center version of the center-surround receptive field. These two areas, if you recall from Chapter 3, work in opposition to each other and pretty much cancel each other out when the entire field is filled. On the left half of the screen is the receptive field with the center and surround regions indicated by the cyan circles on the screen. There are also plus and minus signs indicating the excitatory and inhibitory regions of a stimulus on the receptive field, and the y-axis being the firing rate of the cell at that location. The stimulus you will be using is a white field; say a white sheet of paper against a black background. Currently the stimulus is off the screen represents the background firing rate

for this cell and is provided as a visual reference. In this way, you can compare the cell's response to its response when nothing is stimulating the cell. You can use the slider at the bottom of the screen to move the white field of light across this receptive field from left to right so that eventually it will be entirely covered. As you move this white patch to the right across the receptive field, the firing rate of the cell to the stimulus in that position will be indicated on the graph. By the time you have moved the white patch across the it, you will have a record of how that cell will fire for every position of the edge of the stimulus on the receptive field. In your first pass, simply move the white field across the screen. Also note, that the first time the patch is moved across the screen, you may notice some delay as there is greater computational load the first time. These data are collected and held in the program so that the responses are much faster afterwards. It is important that you move this stimulus across the screen to find the locations of the edge of the stimulus where the cell gives the greatest and lowest firing rate. In particular, notice that the cell gives the lowest firing rate occurs when the white field covers the patch fills the entire central region and the inhibitory surround on one side only.

On this stimulation, you can also change the cell type to an off-center cell and see how these types of cells respond to the exact same type of stimulus. You might consider from the types of cell responses why we have these two types of cells. You can also change the direction of the stimulus and see that these types of cells respond exactly the same way, regardless of the direction of the stimulus.

Now open Experiment 5.x, Mach Bands and Receptive Fields [link to media] and we will see if we can get these receptive fields to generate some sort of output that will indicate Mach Bands to us. When you first open the animation, you will see a row of these center-surround receptive fields overlapping as they would in the eye. All of the receptive fields are drawn on the screen and filled with a transparent version of the same color. Only one row of receptive fields is shown, keeping the illustration onedimensional and relatively simple. On the right half of the screen is a bar graph showing the firing rates of these cells with, a separate bar representing the output of different cell in order from left to right. The colors of the bars will represent the firing rate. As the firing rate goes towards the highest limit plotted, the color will go from blue to teal. The color that the receptive field is drawn in will be the same color as the top of the bar. If the receptive field fires at or near the minimum plotted, it will also be drawn darker.

Pressing the **Stimulus** checkbox will cause an edge like the one that was used in **Experiment 5.x, Lateral Inhibition [link to media]**. A white region is on the left half of the line of receptive fields and a black region on the right. Examine the firing rate of the cells. At this point, I need to give you a little warning about the speed of this figure. Because of all of the computations involved for each of these cells, there may be a bit of a delay, especially on older machines. While the program updates the data for the graph, there will be text displayed on the screen to indicate that figure is in the process of updating. It should not take longer than 5-10 seconds. If you want to speed up the response, you can reduce the resolution of your screen. Returning to the discussion of the figure. The receptive field on the edge of the white region but still inside is has the highest firing rate. Conversely, the receptive field on the edge but still in the dark region has the lowest firing rate. Notice that you get a pattern of firing very similar to what was seen in the limulus eye, and that fits what was observed with the Mach bands. You can then blur the edge by using the **Blurred** check box which became active when you selected to add the stimulus. The edge is now blurred, and you might find the result similar to your results from the Mach Band experiment. You might go back to **Experiment 5.x Lateral Inhibition in Receptive Fields** and see why the results follow this pattern. In other words, why do the receptive fields in these positions give the results they do?

It appears that our receptive fields do have lateral inhibition like the limulus, and that this lateral inhibition is inherent in the organization of the receptive fields. This lateral inhibition enhances the edges, making them more pronounced than they would have been otherwise.

The preceding discussion reveals some important features about how science uncovers ideas about how nature operates. The discussion of the role of contours in perception began with some general observations. Then the discussion moved to studies in a simple system, in this case a simpler organism, the limulus. Next, since the limulus is so different, it was helpful to see if there was any reason to suppose that those findings were relevant to human vision. The observation of Mach Bands was helpful here, suggesting that lateral inhibition may exist in the human visual system. Last, the discussion returned to the discussion of known mammalian physiology of the visual system, including human. These final considerations allowed us to verify that what is known about mammalian receptive fields can create these edge effects.

While the actual pattern of discovery did not follow this path, ideas were draw from numerous sources and tied together. Here there were general observations, work in simpler organisms, work in mammals, and, finally, careful psychophysical work in humans all coming together to generate a richer and more complete understanding. The physiology suggests mechanisms, and the psychophysics provides support that these mechanisms operate in humans. Often it takes many types and sources of evidence to figure out what is going on in a given phenomenon, in this case the importance of edges to how we see. **The Role of Contour in Perception**

So far we have only examined one illustration of the role of lateral inhibition in vision. Let us examine some more of the many different phenomena that the lateral inhibition found in our centersurround receptive fields play a role in. Open **Experiment 5.x: Lateral Inhibition Effects [link to media]**. This screen will be divided into two main regions. The region on the top left will be the stimulus being investigated, with a drawing representing a set of center-surround receptive fields like you saw in **Experiment 5.x, Mach Bands and Receptive Fields**. Below the stimulus with the cells will be the graph which will plot the output of the cells. The bars will approximately line up with the associated cell, to facilitate your understanding of what is being represented. On the right is a panel where you can control the stimuli and cell type. You can select the stimulus that you want to observe, using the **Stimulus Type** drop-down menu at the top of the stimulus panel. For some stimuli, some additional controls will be visible. They will be discussed below with the relevant stimuli. What we will be examining in this experiment, is what the output of the cells can tell us about our perception of these various stimuli. Before we do, since this is the same set of cells as in **Interactive Figure 5.x, Mach Bands and Receptive Fields**, recall that there may be a few seconds delay. During the update period, it will be indicated at the bottom of the graph in the same manner as before.

The experiment opens with the sharp edges version of the Mach Bands stimuli that we have seen before. The name **Mach Bands** appears both over the stimulus in red to contrast with the background, and in the **Stimulus Type** dropdown menu. This example is included for completeness and so that you can see that this model works like the one you have just seen. Therefore, compare the results of this stimulus to what you saw on the last stimulus.

In the drop-down menu, select **White Page** as the stimulus. This is not a phenomenon that has been discussed, but it will help us understand some of the other phenomena that will be discussed. First, observe that in the white page, the response from the center of the output is about the same as the response from the dark region outside of the white page. When a cell's receptive field is filled, it has a very similar firing rate to when it is the completely dark. You might take a while to consider what role this behavior might play in dark/light adaptation. Going back to the stimulus, the main responses of these cells are to be found at the edges. The responses at the bright side of the edge are elevated relative to the middle of the white page, and the responses at the edge of the dark region are depressed relative to the middle of the dark region.

This pattern of firing may be reminiscent of something you have seen before. Go back to the drop-down menu and select the **Craik-** stimulus, which is the Craik-O'Brien-Cornsweet illusion you have seen before. You can also add the luminance profile you saw before by clicking on the Show Profile check box that appeared when you selected this stimulus. It appears below the phrase Other Controls on the right hand side of the screen. Notice that this stimulus is really simple a set of edges. This profile looks a lot like the output to the white sheet, which was a response just at the edges. The output of the cells also is just principally a response at the edges. Here is an important point to take away. While we don't just see edges with either the white page or the Craik-O'Brien-Cornsweet illusion so there is much to be explained, they both have very similar responses at the level of the receptive fields of the eye. Now all the later cells and pathways in our visual system get their signals from these cells, and these later regions of the visual system are what determine our perceptual experience. So, if two stimuli cause the same type of behavior in cells of the eye, regardless of what type of input caused this response, these two stimuli will have to look the same because these later cells get the same input. Here, a white page and the Craik-O'Brien-Cornsweet illusion cause similar effects at the level of the retina, so they appear the same. The process that allows the responses of edges at the level of the retina to be seen as a solid object is called filling-in (REF). Looking more closely at this model, you might not find the particular version of the Craik-O'Brien-Cornsweet illusion shown currently on the screen as compelling as the model seems to suggest. That is because the receptive fields on the screen are much larger than most of yours (REF). To

check that this illusion would work with larger receptive fields, move your face up to the screen where you will use the larger receptive fields of the periphery, and you will see that the illusion works quite well after all.

Ok, we have seen how edges can lead to the perception of an object not there. Return to the stimulus drop-down menu and select **Minimal Contour**. Here we had severely blurred edges leading to the lack of a perception of an object that was there. If you look at the responses of the cells, you will see that they will give the same response, near the background firing rate, to this blurred stimulus. There needs to be a sufficient change in the stimulus within the region of their receptive fields for there to be a good response to the cells.

Finally, let us examine a classic phenomenon. To examine this phenomenon, leave this interactive illustration for a moment and open Interactive Illustration 5.x, Simultaneous Contrast [link to media]. Here we are going to see the influence of the surround on brightness of an object, which is known as simultaneous contrast [to glossary]. The illustration begins with a static view of the phenomenon. The two gray squares are the exactly the same luminance. However, the gray square on the left surrounded by the black region appears brighter than the square on the right surrounded by the bright region. You can make simultaneous contrast even more dramatic by pressing the Start button. When you press this button, the two surrounds will vary over time. The dark side will turn white and then back to black. The light side will do the same thing, but it will be going in the opposite direction. The two center gray squares will not change. Verify this to yourself by covering the surround, and watch that the squares do not change in any way. However, when you pull your hand back you will see both squares changing their brightness a great deal. Each center square will appear to be changing brightness in the opposite direction to its surround.

If you want, you can make an experiment of this little phenomenon. Open Experiment 5.x, <u>Measuring Simultaneous Contrast</u> [link to media]. There are several stimulus parameters that you can adjust in your experiment that show up on the first screen. If you are interested in really exploring this phenomenon you can change these setting systematically across several runs and seeing how they impact your results. For this example just press the **Done** button at the bottom of the screen to proceed to the experiment. This is a simple experiment where you use the stimulus slider to adjust the brightness of the right hand square until it matches the left hand square. When you are done, press the **Match** button, and the right square's surround will be removed, and the two squares will be moved next to each other. In addition, the difference in the computer setting drawing the two squares will be presented. You can also press the **Try Again** button to run the experiment again using the same conditions.

Now reopen Interactive Illustration 5.x: Lateral Inhibition Effects [link to media] and go to the stimulus drop-down menu and select **Simultaneous Contrast**. You will see a medium gray circle with a dark surround. You see the edges on the output of the model, much like you saw with the white page. The highest response is next to the square. Now press the **Increase** button and the surround will increase in luminance and watch what happens to the response of the receptive fields at the edges. When the surround gets to be more intense than the square, the highest response faces the square. From several examples, we know that a filled region will have the brightness determined by the activity of the receptive fields at the edges, regardless of the intensity of the light filling that region. So the brightness of that filling in process must be affected by the direction of the edges that we have been observing. As a result, the object with the highest response gets the brighter perception, and the object with the darker edge will be seen as darker. Look and see if this does not predict simultaneous contrast. You can use the **Decrease** button to reverse what you have done with the surround, and reexamine what you have seen.

Contour and Contrast in Application

There are many ways that the findings discussed above, have important applications. One way is the need to have an image with sufficient contrast on the screen at all times. Take the cathode ray tube (CRT), used in standard televisions, for example. The CRT has a fairly reflective surface that washes out the screen (REF). Usually, you do not want your computer screen or your television near an open window on a sunny day. They can be come hard to see because of the glare. The visual effect of glare is to reduce contrast. As you recall, if there is not enough contrast to support good sharp edges, the object becomes essentially invisible. This problem is bad enough on a computer or television, but consider the situation in an airplane cockpit, where the light from the sun can be much brighter. CRTs have been in commercial

airplane cockpits since the very early 1980's. Before they could be put into the cockpit, however, they needed to be able to be read quickly and accurately. Careful research established the minimal luminance a CRT in the cockpit display needs to generate to maintain a contrast ratio needed to read the display under the worst lighting conditions (Silverstein & Merrifield, 1985; REF).

[NOTE TO ME: Could I find pictures that I can use? Check with Lou about permission]. Another application shows the opposite side of contrast, using reduced contrast edges to hide details that are not desirable. Recall the discussion about jaggies in the last chapter and the use of smoothed lines to hide those annoying steps. To help you remember this discussion open up Interactive Illustration 5.x, Sharp and Blurred Edges [link to media]. This figure is very similar to Interactive Illustration 5.x, Vernier Acuity and Smoothed Lines, but presented in a way to illustrate how using what has been discussed about contrast can improve the quality of image on a computer monitor. When you open the figure, you will see a line drawn vertically. Use the **Tilt Line** slider and tilt the line just a little bit off of vertical. You can easily see the steps in the line, caused by the fact that screen is made up of discrete elements called pixels. This line is draw with just two levels of intensity, full on for the line and full off for the background. Use the **Number of Gray Levels** to allow the routine drawing the line to use some more gray levels, to make these steps have less contrast. As you add more gray levels, these steps caused by the pixels become less visible and what you will see that the line looks like a smooth line drawn without any evidence of the pixels on the screen. [NOT DONE YET: Clicking and dragging the mouse on the screen will move the line around as well, if you wish to see the effect of drawing the line on different screen positions.] It is now apparent, you can see that this hiding of the steps of the line is due to reducing the contrast at the edges of the line. This reduction of contrast at the edges hides the jaggies in the same way as in the minimal contours example (Krantz, 2000; Silverstein, Krantz, Gomer, Yeh, & Monty, 1990). Limits to Lateral Inhibition as an Explanation

Open Interactive Illustration 5.x, Brightness Assimilation [link to media]. In this figure we will see an illustration that shows that not every perceptual phenomenon can be explained by lateral inhibition in center-surround receptive fields. When the screen first comes up, you will see two sets of horizontal bars. The bars on the left are alternating white and gray, while the bars on the right are alternating black and gray. The gray on both sides is identical. Probably the gray does not appear the same. If you need to be convinced that the gray is identical, adjust the **Assimilation Bar Size** to 0, which will remove white and black bars. When you are done return the Assimilation Bar Size to default value, which matches the current **Bar Size**. The gray with the white bars looks lighter than the grav with the black bars. If you recall simultaneous contrast, this effect is opposite of that one. This phenomenon is called **brightness assimilation [to glossary]**. Part of the explanation seems to be related to the widths of the bars, as can be see by adjusting the Assimilation Bar Size to larger sizes. This slider controls the width of the white and black bars, which are called assimilation bars here because the gray assimilates to these brightnesses. As you make these assimilation bars wider, the effect will go from brightness assimilation to simultaneous contrast. You can also examine the width of the gray bars (Bar **Size**) and the gap between the two gratings (**Gap Size**) to see how this influences your perception. Brightness assimilation shows that the explanation of form perception based solely on retinal receptive fields is incomplete. It is not so much wrong, as there are other levels of the visual system that will have their role in how form perception works. This distinction between a theory being wrong and incomplete will be seen in other features of our perceptual system, most notably color vision (Chapter 6) and pitch perception (Chapter 10). For now, it is now time to examine the role of some of the other portions of the visual system on form perception (Shapley & Reid, 1985).

b. Feature Extraction [I am not sure what I want to do here] – Hubel and Wiesel?

Patterns across Regions

There is certainly more to a stimulus than its edges. Most stimuli are far more complex than the simple white fields that we have been examining. For example, take a face. The light luminance changes throughout most of the face, and edges are simply areas where the change in the luminance is greater. To be able to see a face, our perceptual systems need some way to be able to extract these changes and how they proceed across the face. There have been several approaches, and many of them have used an analogy drawn from a mathematical theorem known as Fourier's Theorem. Therefore, we will start with a description of this theorem so we can understand some of the ideas. It will be useful to try to understand

the basic ideas of this theorem, as it will play a roll in the discussion of audition as well as vision. Also, recall that science use analogies from well understood ideas or phenomena to explain less well understood phenomena. This is the case here. No one is going to claim that the visual system does what will be called a **Fourier Analysis [to glossary]**, but there is something going on in the visual and auditory systems that seems similar. However to feel that time is well spent understanding this theory, it is useful to give some of the background concepts that led to a consideration of this theory as an analogy.

Spatial Frequency and Sine Wave Gratings

As was stated in the first chapter, science likes to keep its experiments as simple as possible to help clarify what is going on in nature. The experiment is very simple with only the independent variables changing; and these variables change only in a very controlled fashion. In addition, the stimuli, as you probably have noticed by now, are often very simple. That way it is clear what is being manipulated in the stimulus and it is clearer what the visual system is responding to. Using a complex scene that has so much going on in it, for example, a picture of a river, it might not be clear what ought to be manipulated, or even how changing one feature of the stimulus, say the contrast, will affect the rest of the stimulus. However, as we shall see below, it is not always clear what makes a stimulus simple to the visual system. Therefore, the search for the best stimulus to investigate the visual system is an ongoing one.

In the search for different stimuli that are simple to the visual system, one of the stimuli that has been used is the sine wave grating. Open Interactive Illustration 5.x, Types of Gratings [link to media]. In this illustration, we will explore the nature of sine wave gratings more thoroughly as well as how to characterize them. When you open the figure, there are two different gratings being displayed. On the top of the gratings is a plot of the luminance for that grating. On the left is the grating of the type that was discussed in the section on acuity. This type of grating is often called a square wave grating. The square in square wave grating because its luminance varies according to the sine function. Both gratings have bars of the same width and the same contrast ratio. The sine wave grating looks like a blurry version of the square wave grating, and as you shall see later, that is not a bad description.

The sine function used for the sine wave gratings is also the same sine function we used to illustrate the waves for light. There are two measures of this sine function, frequency and intensity. The frequency is how many bars there are in a given region, and it is inversely related to wavelength. At the bottom of the screen is a slider you can use to change the frequency of both gratings. As you move the **Frequency** slider to the right, the frequency of the gratings will increase, and you will see that the bars get smaller. On the luminance plots you can see how the distances between the peaks get smaller. Below the **Frequency** slider is the **Contrast** slider, and contrast is related to the intensity of the waves and, thus, the gratings. As you adjust the **Contrast** slider, the vertical distance between peaks and troughs in the waves get smaller and the gratings get less distinct, until eventually it is impossible to see them.

For several reasons, some of which will be discussed below, researchers have used the sine wave grating a great deal as a research stimulus. The need, then, is to find a way to develop good measures to characterize these gratings. The first is related to the intensity and the measure is contrast. In our case we would use the contrast ratio between the luminance at the peak and the luminance at the trough. The other measure needs to be related to the frequency of the grating. Since this frequency is spread out in space, it is called **spatial frequency [to glossary]**. From acuity, we know that the size of an object on the retina depends upon how far away the object is from the eye, and that it is the size of the object at the retina that is important (refer to Interactive Illustration 4.x, Visual Angle if you need a refresher). So the measure of the spatial frequency needs to be related to visual angle. The standard unit is the number of cycles that fit in a degree and visual angle: cycles/degree for short. The smaller the bars, the higher the frequency, and the more will fit into the degree, so the higher the number the higher the frequency just as in the illustration.

The Contrast Sensitivity Function

With these measures of the sine wave grating, there appears a new threshold that can be measured. This threshold is: for a given spatial frequency, the smallest amount of contrast necessary to reliably make out the grating. What is interesting is that the contrast threshold depends upon the spatial frequency, as will be illustrated in the following experiment: **Experiment 5.x: Contrast Thresholds and Sensitivity [link to media]**. When you open this experiment, you will be presented with a rather odd stimulus. There is a grating in the upper left hand portion of the screen, but it is not a simple sine wave grating. As you look from left to right, the spatial frequency increases regularly. As you look from bottom to top, the contrast

goes from relatively high to zero. To do this experiment, use the following instructions: First, find a comfortable distance and sitting position. You need keep your head is as constant a position as possible, because moving your head will change the spatial frequencies you are viewing, and we want them to stay as constant as possible. Normally experimenters would run their subjects in a head and chin rest for this type of experiment to fix their head position, but you will have to keep your head relatively still without this device. At the bottom of the grating stimulus are short vertical cyan lines. These lines are the regions where you are going to make your judgments. In the space between two of these lines, find the highest point on the bar that you can see in that region and click on it. You will get a red dot as feedback to your click. The higher the position of the point on the grating, the lower the threshold. There are 10 regions, getting smaller to the right in the grating. You will make this contrast judgment, the new click will replace the old. You can also go back to any region and refine your judgments. Make sure you do not move your head too much, as that will change your results in important ways. It does not matter too much what distance you are from the screen, but keeping still is important. Now, do the study and then return to the text.

You probably found that your results form a slightly off balance upside-down "U" shape. We are most sensitive, meaning we can detect the smallest contrast, to the spatial frequencies near the middle of the figure. If you just glance at the grating without the data, you will see this pattern as a whole. You can clear the points on the screen by clicking the **Hide Points** button. It will not lose the data. You can restore the plot of the data on the grating by pressing the same button again, which now says **Show Points**. You can clear the data by clicking **Clear Data** button. This button allows you to repeat the study. Pressing the **Show Data** button will allow you to retrieve the data. The spatial frequency is in cycles per the grating, and the contrast is an approximate value, but it can give you some idea of your the contrast sensitivity function.

The typical way to present this data is not as thresholds but as **sensitivity [to glossary]**. Sensitivity is simply the inverse of the threshold, so where a threshold is low, the sensitivity is high. Your data is actually a measure of sensitivity, since the lower your threshold, the higher up the graph it is plotted.

This curve that plots contrast sensitivity as a function of spatial frequency is called the **Contrast Sensitivity Function [to glossary]**. Notice that we are most sensitive to spatial frequencies in the middle range. These values are usually given in the range 4 to 6 cycles/degree (REF). We are less sensitive to higher and lower spatial frequencies. The higher spatial frequencies are the smaller bars and represent smaller objects. At some point, there will be a spatial frequency that we cannot see, regardless of the contrast. This point will be the same as our acuity limit. So, if we examine a person's contrast sensitivity function, we learn a lot more than from simply determining their acuity. The person's acuity can on the contrast sensitivity function. It is the highest spatial frequency that can be resolved at even the highest grating. But, there is a lot more information about a person's ability to see than just this limit. However, what to make of all of this information ,and what it may mean for our perception depends upon what can be learned from the concept of Fourier analysis.

Fourier Analysis

Fourier was a 19th century French mathematician who discovered a simple way to describe complex patterns. In particular, the patterns that were being investigated by Fourier were waveforms. You might immediately assume that the waveforms in vision that we might be interested in would be the waves associated with the individual wavelengths of light. The waves of light, however, are not the feature of vision being used in the analogy to the Fourier analysis. Open <u>Interactive Illustration 5.x</u>, <u>Waveforms in</u> <u>Objects [link to media]</u>, and it will illustrate where these waves that will play a role form perception arise. In this illustration, there is a picture on the left half of the screen and a graph on the right side. Press the **Start** button at the bottom of the screen and it will put a red horizontal line across the picture. The graph on the right will show a plot of the luminance across the picture at the level of the red line. Moving the **Level in Picture** slider at the bottom of the screen will move that line up and down in the picture. You can also click in or drag your mouse in the picture to move the line in the picture as well. The plot on the right will keep up and indicate the luminance at that point in the picture. You can use the **Picture** drop down menu to pick a new picture to view. There are 4 pictures that can be selected. Notice how, regardless of the picture chosen or the location of the red line in the picture, the plot on the graph is a line that goes up and down irregularly. This is a complex wave, and to the scientist it begs to be simplified in

some way to be described more easily. Also notice how the pattern is different both for different locations in the picture, and for the different pictures. The idea is that there is some information in these differences of this complex waveform that the eve can take advantage of to see objects in the world.

[ADD A DISCUSSION OF THE ADDITION OF WAVES ABOUT HERE]

That is what the Fourier analysis aspect of Fourier's Theorem is all about. It describes a way to take these complex waves and describe them in terms of a much simpler set of waveS and then these simpler waves can be examined for patterns hidden in the complex wave. Usually, sine waves are used as the simpler waves. There is no need to deal with the mathematics (I think I heard that relieved sigh here at my computer), but using a visual illustration of the main concepts of how this works will be very helpful. Open Interactive Illustration 5.x, Fourier Analysis [link to media]. In this illustration we will take one complex grating and see what happens to that grating after Fourier analysis. This example will allow us to get some feel for how Fourier analysis describes an image.

When you first open the screen, you will see in the top left quadrant of the screen what in Fourier terms is a complex image: a bar grating. The light and dark bars do not resemble the sine waves that will be the simple waves used in the Fourier analysis. On the right side of the screen is the output of the Fourier analysis of this square wave grating. This is a plot that shows which sine waves are needed to generate the square wave grating. The x-axis of the graph is the relative frequency of the sine wave. Recall that the frequency of a wave is the inverse of the wavelength. The higher the frequency, the shorter the wavelength. The lowest frequency is called the fundamental and is labeled 1 in this illustration. It is a sine wave, as we shall soon see, with the same frequency as the bars in the square-wave grating. The other frequencies have frequencies that are whole multiples of the frequency of the fundamental. So the next bar is a given a 3, so there will be three times as many cycles of that frequency in the same area as for the fundamental. The heights of the bars refer to the amplitude of the sine wave in this grating. Actually, this graph only represents the first few sine waves that come out of the Fourier analysis of a square wave grate. To make the grating perfectly square requires an infinite series of sine waves.

Before we go on with the figure, let us consider the grating and the graph a bit and what they mean. The grating and the graph are actually two different representations of the grating. The grating is the way the grating appears laid out in space and is referred to at the **Spatial Domain**, as shown on the figure. The graph is the same stimulus, but plotted with the x-axis showing the frequency of the sine waves needed to make this grating in the spatial domain. So this version of the grating is called the **Frequency**

Domain.

The bottom half of the screen is in some ways the same as the top half but done in a way to allow you to manipulate the sine waves in the grating, so you can see the impact that these sine waves make in this square wave grating. On the left is the sum of all of the gratings shown in the graph above. You have something that is near a square wave grating, but not quite. Remember, this graph only has the first few elements in the square wave grating. The sliders on the right bottom portion of the screen control the amplitude of all of the sine waves in the grating on the left. In addition to the slider, there are several other controls that need to be described. First, the **Zero All** button will make all of the gratings have a zero amplitude, and wipe out the grating on the left. Try it. To get back to the beginning situation, press the **Reset All** button. Below these two buttons are two check boxes, **Grating** and **Wave**, which control what is displayed on the left side of the screen. Currently the **Grating** is checked, which means the gratings are being displayed. If you check the **Wave** check box, an orange line will be drawn over the grating to show how the luminance changes over the screen surface. It will be perfectly in step with the grating and may help you see how this Fourier analysis works. Below each slider are three small buttons. The functions of these buttons are labeled to the left of the buttons. **Show Alone**, the top of the three buttons, will cause only the sine wave of that frequency to be displayed in the grating to the left. **Reset**, the middle of the three buttons, will restore the amplitude of that sine wave to what is needed for a sine wave grating. **Zero**, the bottom of the three buttons, sets the amplitude of that sine wave to zero, effectively removing it from the grating on the left. It is time now explore Fourier analysis using this illustration.

Use the **Show Alone** button below the slider for the fundamental (the slider labeled **1**). This will take all of the other sine waves and turn them to zero, leaving the fundamental. Now the graph and the image on the bottom line will be of a sine wave and sine wave grating, respectively. This grating has two

cycles across the grating which are the number of fuzzy dark and light bars that you see on the screen. This is exactly the same number of cycles as in the square wave and square wave grating at the top of the image. There is something important to notice right here. The fundamental tells us a lot about the general shape of the grating. Now press the **Show Alone** button under the third harmonic, that is, the sine wave with three times the frequency of the fundamental and it is the slider labeled **3**. Count the number of peaks on the screen. You will see three times as many as you did with the fundamental and three times as many as there are bars in the square wave grating. It might seem odd that this frequency would be of any use to the perception of a square-wave grating but actually it is very useful. Press the **Zero All** button to clear the grating and then press the **Reset** button under the **1** slider first. Notice the shape of the grating. While watching the grating, press the **Reset** button under the **3** slider. The **Reset** button under any of the sliders returns that slider to the value it was as a result of the Fourier Analysis. The bottom grating now shows what an image looks like that has these two sine wave components. Notice a couple of things: first, the amplitude of the third harmonic is not as great as the amplitude of the fundamental. In fact, it turns out that the amplitude of the third harmonic is $1/3^{rd}$ that of the fundamental. Not particularly important, but an interesting coincidence. In the graph, click the **Zero** button and then the **Reset** button below the **3** slider to remove and return the third harmonic, and watch what happens to both the graph and the image on the bottom row. Notice that the combination of the fundamental and 3rd harmonic has steeper sides and a flatter (if bumpy) top than the fundamental alone. In other words, the combination is more like a square wave than the individual components. With the fundamental and third harmonics reset, add successively higher harmonics. You need only to add those harmonics that have odd numbers. Only those sine waves that are an odd multiple of the fundamental are in a square wave. So keep adding more and more harmonics and watch what happens to the combined figure. As you add each harmonic, add and remove it using its **Reset** and **Zero** buttons as you did with the third harmonic. These dynamic changes to the graph will make what is happening easier to see. When you have finished adding and removing all of the harmonics, come back to the text.

As you finally added all of harmonics with ever higher frequencies, the graph became more and more like the square wave grating. The edges become sharper, and the top becomes flatter. It is never identical to a square wave grating, because we really do not have enough sine waves to make a complete square wave, but eventually the grating would get so fine that it is above our acuity threshold, and the combination of all of the sine waves would appear identical to the square wave.

You might look at all of this mess and see the square wave as a lot simpler than the sine waves. Well, there are two points to make. First, recall that this is an analogy. No one will argue that the visual system does a Fourier analysis. Just that there is something in the visual system that works this way. We will discuss some of these ideas next. Second, yes this output is more complex than a simple bar graph. But recall that the world is a very complex place. Many different types of shapes and patterns exist. If we had to have something that responded to all of these shapes and patterns, it would be very complex. However, in Fourier analysis, the same sine waves can be used to see all of the shapes and patterns. Play with the illustration, and see how having these few sine waves can generate a lot of different outputs in that figure.

Just to show that this situation with the square wave grating is not made up, you can use the **Grating Type** drop down menu to select a **Triangle** grating where the luminance changes linearly to points. Using the **Wave** check box to show the waveform is very useful in this type of grating. This grating is even better matched by the few sine waves that this illustration has available.

Bringing the Contrast Sensitivity Function and Fourier Analysis Together

The pieces are now in place to understand how the contrast sensitivity function gives us a much more complete picture of the way that we see objects. The use of Fourier analysis as an analogy suggests a possibility for the functioning of the visual system. The possibility is that one way for the visual system to represent the myriad forms we see would be to have detectors for spatial frequency in the brain. That is, some part of the visual system would respond to the fundamental spatial frequency or general shape of an object. Other parts of the visual system would respond to the highest spatial frequencies that give the details and sharp edges, as in the square wave grating shown in **Interactive Illustration 5.x**, **Fourier Analysis [link to media]**. The advantage of this arrangement is the same advantage as was found in Fourier analysis, one set of response elements would work for all objects. We would not need one set of

detectors for round objects and different set for square objects, and even more different sets of sets for all of the irregular objects that are more common in the natural world than in our constructed world.

The Concept of Channels. So far the discussion has been very abstract, so let us try to make this discussion more concrete. The elements of the visual systems would be cells at some level or more. These cells would be constructed to respond best to different sets of spatial frequencies. There does not need to be a different cell for different spatial frequencies. In fact, it is possible to use a very small set of cells and still give the brain a response that indicates exactly the spatial frequencies being viewed. To cast a look both backwards and forwards in this text, this same ability to use a very few cells to recover a lot of information is what is going on in color vision that allows us to experience many different colors with only three classes of cones. These different parts of the visual system that are sensitive to a range of stimuli are called channels.

Open Interactive Illustration 5.x, Spatial Frequency Channels [link to media NOT DONE]. In this activity, you will get a chance to see how a few channels sensitive to spatial frequency can recover the presence of many more spatial frequencies than there are channels. Pay close attention, as this same illustration will be slightly modified and used to explain how we see color. Learn it here and save yourself a lot of work later. When you open the illustration, the screen will show two sine wave gratings on the left half of the screen. The top grating is labeled the **Standard**, using the terminology from psychophysics for the stimulus that is to be matched. The bottom grating is labeled the **Comparison** and this grating you will try and match to the standard. However, you will not match it in your visual system. You are going to use the spatial frequency channel indicated on the right half of the screen. The sensitivity of the channel to different spatial frequencies is plotted at the top part of the right side of the screen. This plot is analogous to the spectral sensitivity function you saw in Chapter 4 for the rods. However, instead of this being the sensitivity of a cell to different wavelengths of light, this is the sensitivity of a collection of cells to different spatial frequencies. This channel is most sensitive to some spatial frequencies, and much less sensitive to others. The response of this channel will depend both on the spatial frequency and the contrast of this spatial frequency. The firing rate of the channel to the **Standard** grating is plotted with the bar on the left of the graph, below the plot of the channel. The firing rate channel to the **Comparison** grating is plotted on the right hand bar in the same plot. You task is to make these firing rates match each other. I am going to make it difficult for you. You can only adjust the contrast of the **Comparison** grating. You adjust the contrast of the **Comparison** grating by using the **Contrast** slider to the right of it. When you have a match, the word **Match** will appear on the screen. Make this match and then return to the text.

First, notice that it was not terribly hard to make this match. Sliding the contrast of the comparison up and down, it was rather easy to find a point where the firing rate of the spatial frequency channel to the two gratings is the same. Second, notice that the two gratings do not look at all the same to you. Before we pursue this second point, let us further elaborate the first point. Click the **New**

Comparison button below the comparison grating, and a new spatial frequency grating will be in the comparison window. Again see if you can make a match in this spatial frequency channel. Try it a few times and see if you can find a grating that you cannot match. You will find that no matter what spatial frequency the comparison grating has, you can match it to the standard. You can press the **New**

Standard and try it again for a different standard and get the same results. Regardless of the spatial frequency of the standard and comparison, it is possible, even rather easy, to find two contrasts for the two gratings where they cause the same firing rate for the channel. Well, if the channel is firing the same way in the two cases, then the brain should see the two gratings exactly the same way. All the brain has is the firing rate of the channels, and these firing rates are the same in the two cases. Now go back to the second point, that these two gratings do not look at all the same to you when you make these matches. So it seems likely that we are not built this way.

Press the **Multiple Channels** button below the graph of the channel. Three different spatial frequency channels are present instead of one, and each one is drawn in a different color. The graph of the firing rates has three bars for each stimulus; the bars are different colors and match up with their respective channel and represent the firing rate for that channel. Now to make two stimuli look identical it is necessary to have the firing rates of all three channels be exactly the same. If they are not the same, then the brain will have a difference in at least one of the channels that it can use to see that there are two different gratings. Let us see how this works. The rest of the interactive illustration works the same way as

before. Use the **Contrast** slider beside the comparison and see if you can adjust the contrast to find a match. Move it slowly and carefully, and be sure of your conclusion. Once you are sure if you can or cannot make a match, then press the **New Comparison** button and try another comparison. You can also try different standards to make sure it is not some fluke of the particular standard picked that is causing your results. Try some examples and return to the text.

These results are very different from the single channel results. Whenever the spatial frequency of the **Standard** and the **Comparison** are different, it is impossible to make a match. At least one of the channels has a different firing rate, no matter what contrast is used. If you press the **Same SF** button below the **Comparison** grating, which became visible when you went to multiple channels, you can see that you will only get a match between the **Standard** and **Comparison** at one contrast as well, and it is when the contrast between the two gratings are the same. In this case, a match is only possible when the two gratings appear the same to you.

There are again two things to notice in this multiple channels model of seeing spatial frequencies. First, each spatial frequency can be discriminated as being different by the presence of multiple channels. Second, the match in the channels occurs only when there is a match perceptually. This is not to say we have three spatial frequency channels. The claim is much less bold than that. What is being claimed is that we might have something like these spatial frequency channels in our visual system.

While all of this discussion is nice on a theoretical level, what is needed is some evidence that these channels that are sensitive to different bands of spatial frequency actually exist. While there have been many studies that have supported the idea of spatial frequency channels in a general sense, one phenomenon has received the most attention, and the discussion of this phenomenon will serve our purposes here.

The Blakemore-Sutton Effect

The Blakemore-Sutton effect first published in 1969, is a classic study that will allow us to understand better the concept of spatial frequency channels as well as an important psychophysical study technique, adaptation. We will take up each of these two concepts first, and then do an experiment and interpret the results of that experiment.

Adaptation and Aftereffects as a Research Tool. Open Interactive Illustration 5.x, Adaptation and Aftereffects [link to media NOT DONE]. In this illustration, you will experience a very basic example of a phenomenon that you probably have seen fancier examples of already. We will use this simple example to discuss the significance of these effects as a research tool. Press the **Start** button, and then stare at the center of the white square. After about 45 seconds the square will be replaced by a white screen. Note what you see when the screen goes entirely white and then return to the book.

Most people report that they saw a dark square where the white square had been. This is an **aftereffect [to glossary]**. You were staring at the bright, white square for a long time, and then it was removed. At this point, you see the opposite perception, in this case a dark square. In some sense, you are less sensitive to the light that falls on the same region of the retina where the square had been. This process of staring at a stimulus for a long time and this staring altering perception so that an aftereffect occurs is called **adaptation [to glossary]**.

The fact that it seems that adaptation leads to a loss in sensitivity to the stimulus adapted to, led some researchers to claim that adaptation fatigues or tires those channels sensitive to the stimulus. In this case, the sensitivity to the bright stimuli was reduced. So adaptation, when carefully done (REF) can be a method of teasing these channels out of the visual system. By adapting part of the visual system and seeing what aftereffect occurs, then there is a chance to understand what channels may exist in the visual system.

Spatial Frequency Channels and Adaptation. So if I take a grating and adapt a person to it, then what will the aftereffect look like if we have multiple spatial frequency channels? That is the question that faced Blakemore and Sutton as they developed their now famous aftereffect. Exactly what the aftereffect will appear to be, will depend upon what stimulus you use after the adaptation. In the example above, a white screen was used. What do you think would have happened if you just put a person's face there after the adaptation to the white square? Blakemore and Sutton proposed using a sine wave grating, one that had a different spatial frequency from the sine wave grating that was used in adaptation. Open **Interactive Illustration 5.x, Adapting Spatial Frequency Channels [link to media NOT DONE]**. In this illustration, you will be able to theoretically adapt a grating and see what happens to the appearance of

other gratings. This illustration will help you develop predictions for the outcome of the Blakemore-Sutton experiment you will do next.

When you open the screen, you will see a plot with three spatial frequency channels, just like **Interactive Illustration 5.x, Spatial Frequency Channels [link to media]**. To the left will be a grating, and the vertical black line on the graph of the channels will represent the spatial frequency of this grating. The bar graph below the plot of the spatial frequency channels indicates the firing rate of the three of them. Again, the color of the bar matches the color the corresponding spatial frequency channel. These bars are proportional to the height of the curve where the line for the grating crosses them.

Press the **Adapt** button above the grating. There are two gratings visible and three vertical lines on the graph. The red line is the spatial frequency of the grating that you have stared at for a long time. You have adapted to that grating. Using the idea of fatigue, each of the channels that responded to that grating is now less sensitive to all gratings that they respond to. How much less sensitive the channel has become depends upon how strongly that channel it responded to the adaptation grating. In this example, the spatial frequency channel drawn in blue is the most, fatigued because it responded most strongly to the adapting grating. The adaptation is drawn by making the curve lower, with the original curve drawn in a darker blue, and similar changes are made for the other channels. The black line on the graph is still for the original grating, but that grating is labeled the **Unadapted or Input** grating.

It is important to examine the responses of these channels to the new grating. Examine the bar graph. The original firing rate across the three gratings is shown, and the firing rates are shown for the grating that it is adapted. Notice that the firing rates change, and change in terms of how strongly the different channels fire. The brain only knows firing rates, and these firing rates do not tell the brain that they are from an adapted channel. That is very important to realize. Since the firing rates change, then the perception must change. That is what the green line and the **Adapted** grating below the **Unadapted or Input** grating refer to. This grating has the same pattern of responses on the three channels as would be found for a grating if there was no adaptation. Notice that the adapted grating has a higher spatial frequency than the unadapted grating. That is because the unadapted grating has a higher spatial frequency than the grating that you adapted to. If you tested with a grating that has a lower spatial frequency than the one you adapted to, the grating would appear to have an even lower spatial frequency than it does. In general, after adaptation, gratings appear to have spatial frequencies that are farther from the adaptation grating than they actually are.

<u>A Blakemore-Sutton Effect Experiment.</u> In this experiment, you will get the chance to see how an adaptation experiment can be used to determine some of the ways that the visual system operates. In this case, you will adapt to a sine wave grating and then see how it alters your perception of the size of another one. Open <u>Experiment 5.x, Blakemore-Sutton Effect [link to media]</u>. The first screen is the screen where you can setup the experiment or demonstration. The default values allow you to demonstrate the effect first, and you can try that demonstration by pressing the **Done** button at the very bottom of the screen. When you press the **Done** button, you will see a gray screen with the **Start** button at the bottom of the screen. When you press that button, you will see two gratings with a white fixation bar between them. You look at the fixation bar, moving your eyes from one end to the other at a slow regular pace, say moving your eyes once a second. This is the adaptation phase. The bottom part of your retina is adapting to the spatial frequency of the grating above the fixation mark. This is a fairly high spatial frequency grating. Your retina above the fovea is adapting to the spatial frequency of the lower grating, which is of a relatively low spatial frequency. After 60 seconds, these two gratings will be replaced by a two identical gratings with a spatial frequency in between the two adaptation gratings. This is the testing phase. The question is, do these two gratings look the same? Try the demonstration and see what you find.

Since the time of Blakemore and Sutton's first paper, people have seen the top grating during the testing phase as having larger bars than the lower grating which is exactly what was predicted from the idea that we have spatial frequency channels. You can repeat the experiment by pressing the **Start** button again. If you are interested, you can test a lot of issues related to this phenomenon by reentering the illustration and running it in the experiment mode when your first enter the illustration. First, the whole process runs a little differently if you select **Experiment** from the **Type of Test to Perform** drop down menu. After the **Stimulus Parameters** window, you will get a **Psychophysical Setup** window for method of adjustment. You will be able to select the number of trials, and how much the

comparison will vary from trial to trial. You will still press the **Start** button to begin the adaptation. During the adaptation phase, you will only adapt to the top grating. During the test phase, the bottom grating will be the unchanging standard stimulus and you will adjust the spatial frequency of the top grating which is the comparison stimulus, to match the spatial frequency of the bottom grating. When the test phase happens, you use the **Adjust Comparison** slider at the bottom of the screen. You can also use the up and down arrows and the page up and page down buttons. When the two gratings look to have the same spatial frequency, press the **Done** button at the bottom of the screen. After you have completed all of your trials you will be show the adaptation grating spatial frequency and that standard spatial frequency again along with the mean and standard deviation of the comparison grating spatial frequencies you thought matched the standard grating.

Some of the parameters you can adjust for the experiment are: the spatial frequency of the adaptation grating, the duration of adaptation, the spatial frequency of the standard, and the type of grating, either square wave or sine wave. I would recommend using sine wave gratings for experiments. Recall all the spatial frequencies present in the square wave grating. However, the square wave grating works nicely for demonstrations. You can also change the size of the region of the screen that has gratings, and if the gratings are tilted during adaptation and/or testing.

If you look at the slider for the **Adaptation Grating Spatial Frequency**, you will see that the units used are the number of cycles that would cover the screen, if the whole screen was visible. This type of unit is used because I don't know the size of your screen or how far away from the screen you are sitting. You can convert these units to cycles per degree, using the calculator at the bottom of the screen. To use this calculator, enter the number of cycles per screen in the top box where it asks **What is the width of your grating**. Measure how far your eyes are from the screen where it says **Viewing Distance**. And then enter the number of cycles per screen where it wants to know the **Number of cycles in the grating**. Then press the **Calculate** button and it will tell you the cycles per degree. If you do not use the calculator during the experiment, you can get to it in the <u>Calculator, Spatial</u> <u>Fequency Caclulator [link to media]</u> which can be found right below the link to start the Blakemore-Sutton experiment on the media web page.

Using these variables, you might examine if the size of the aftereffect is the same for all spatial frequencies of the comparison. How does the spatial frequency of the standard alter the results? What about the duration of the adaptation?

[come up with some more questions to direct examinations?]

Physiological Evidence for Spatial Frequency Channels

[dog cells - reinterpret of Simple cells - get some references and write this tomorrow.]

Applications of the Contrast Sensitivity Function

[aging?]

Grouping and Gestalt Contributions to Form Perception

Background

So far the discussion of how we perceive the form of an object has suggested that all we need to know about form perception can be gotten by looking at how the parts of an object are perceived and then building up to the perception of the whole. The stimulus has been broken down in a lot of different ways: into edges, feature, and even spatial frequency components. But there seems to be more to the perception of an object than could ever be understood in terms of these analyses, or at least that was the position taken by the psychologists that formed the Gestalt Psychology movement.

In this section we will examine some of the laws of perception that they discovered and try to better understand the basic issues raised by their movement.

Figure-Ground Perception

Open up **Interactive Illustration 5.x, Figure-Ground [link to media NOT DONE]** and we will use this figure to illustrate one of the basic features of perception, according to Gestalt Psychology. When you first open this picture, you will see a series of black blotches. If you look for awhile you might make out a form. You can press the **Hint** button and it will give you some assistance in discovering the object.

The **Show** button will reveal the actual object hidden in the pattern. Use this button only if you really struggle to find the object. Close and re-open the window, and see if it is not easier to find the object that you know what you are looking for.

This illustration is a great example of some of the fundamental aspects of perception as uncovered by the Gestalt psychologists. First, we do not just see blotches, spots or disembodied features. We see a **figure [to glossary]** against a background, called the **ground [to glossary]**. Once you see the object in the scene, you see this object. All of the rest of the blotches and the picture, and the computer and all the rest that you see serve as a background to this object. In other words, as we examine the world, it is divided into two regions: the figure, which is the object we are examining; and the ground, which is the background to this figure.

There are a lot of differences between the figure and the ground perceptually. Let us use a classic figure, the Kanizsa Triangle to illustrate some of these differences between the figure and ground (Kanizsa, 1976). Open Interactive Illustration 5.x, Illusory Contours [link to media]. When the figure comes up, you will see what appears to be a triangle in the middle of the screen. However, upon closer inspection you will see that actually, what you think is a triangle is an illusion. No triangle exists. The perception of the triangle is created by the three circles at the corners and the lines that are broken by what appears to be the body of this triangle. The black on the screen is the same everywhere but the region that is inside this illusory triangle appears different from the rest. The fact that we see this shape is a strong indication of the power of our tendency to see the world as a figure against a ground. We have segregated two regions that are the same physically and seen differently. These differences from each other are suggestive nature of figures and grounds.

First, the figure tends to appear in front of the ground. Here the triangle looks to be in front of the circles, which is why parts of the circles appear to be obscured by the triangle you perceive, and the rest of the background appears to continue behind this triangle. There is also a tendency to see the figure as brighter than the ground. Examine the triangle that appears on the screen and see if it does not appear a little bit brighter than the background.

[NOTE TO ME: What are some other differences I want to discuss here].

Just as an aside, take your time and manipulate this figure. What characteristics of the stimulus pattern are necessary to see the figure? What characteristics make the figure more vivid and clear? As you go through this section on the Gestalt laws, see if you see these laws working in this figure.

Grouping and Organization

The Kanizsa triangle suggested that to see a figure, it is often necessary to group distinct elements of the stimulus world into a single organized perception. The Gestalt Psychologist developed a set of laws to illustrate the ways that these groupings happen.

Laws of Proximity and Similarity. Open Interactive Illustration 5.x, Proximity and Similarity [link to media], which will allow us to run through the first two Gestalt Laws. The first law that will be discussed is the Law of Proximity. On the screen, you will see a set of circles. If you move the Horizontal Distance slider at the top or the Vertical Distance slider at the left side of the screen, you can spread out the columns and rows. If you move just one of the sliders, say the slider at the top of the screen, the objects will separate and become visible as separate columns. The pattern on the screen goes from a large number of separate objects to few columns. The grouping is determined by the relative distances of the different objects. Within a column, all of the objects are closer together. Within the rows, which do exist but are hard to see, the objects are relatively far apart. This difference in distance is what drives the perception of the grouping. The columns become grouped as wholes, wholes that are more than the parts. In this case, the grouping is determined by relative distance, and this is the law of proximity [to glossary]. You can adjust the distances between the different rows and columns, and see how the distances will alter the perception of the grouping. For example, ask yourself, how much does the relative distance between the rows and columns have to be to see the columns as columns?

When you have finished playing around with the proximity, move the sliders back so that the objects are as close to each other as possible. Now use the **Shape** slider just to the right to change the shape of the objects that make up the odd numbered columns. Notice that, even though you have the same distance between the objects in the rows and the columns, you still get the impression of columns. You can move the **Shape** slider back, and then move the **Color** slider and see that color can also be used to give the impression of, in this case, rows. Here the grouping in your perception is driven by the similarity

between objects. This is the **law of similarity [to glossary].** The objects that are similar to each other are grouped together perceptually. You can use this illustration to see how the Laws of Proximity and Similarity interact. You can make the columns by similarity, and then adjust the relative distance using the vertical slider on the left side to make rows. Does similarity or proximity seem to have more influence in what you perceive? Does it depend on the amount of similarity and/or proximity? Is shape or color similarity better able to create grouping? What happens when the perception of rows is balanced against the perception of columns? Using the **Size** slider can put more or fewer of the circles or squares on the screen and make your observations easier.

Law of Good Continuation. The law of good continuation [to glossary] deals with how we group contours with each other. If lines cross each other, we have no problem deciding where the line goes after it crosses, in most cases. Open Interactive Illustration 4.x, Law of Good Continuation [link to media]. In this interactive illustration, you will be able to manipulate characteristics of lines and see how changing these characteristics will influence your perception of the stimulus. When you first view the image, you will see what appears to be two curves crossing each other. If you move the Good

Continuation slider at the bottom or the right side of the screen, you will move one of the lines in a fashion that supports this perception. But the lines are really made up of just a pattern of pixels turned on and off of the screen. There is no reason that these lines should be continuous. Is it not also possible that the two lines break where they cross and you have two lines each with a sharp angle in them? Move the

No Continuation slider at the bottom or the right side of the screen, and the no continuation version of the reality will be made apparent to you. Now the lines do not cross but meet at the intersection. However, you have a strong tendency to see the lines as moving smoothly, and when you put the lines back together you will tend to see the smooth version, not the one with the sharp angle. This type of grouping seems to imply that there is a strong bias in our visual system to see lines and contours as smoothly changing unless forced to see the stimulus as having a sharp angle. One instance of being forced to see the sharp angle is to move either of the **No Continuation** sliders where the lines break apart.

Continuation does not require solid lines either. Click on the **Dashed Lines** check box in the lower right hand corner of the screen. This causes the lines to be made up of small dashes, and the perception of continuation does not change in any way. The proximity of the dashes is not involved in why you see continuation in this case, because the dashes for the sharp angle are as close as the dashes for the smooth, good continuation angle.

Another feature of this law you might want to explore the smoothness of the lines themselves where they cross each other. If you adjust the **Angle** slider that you can find on the right side of the screen, you will bend one of lines at an angle at the intersection. See how this influences your perception of the continuation. Another question you might ask is how this law interacts with the law of similarity. You can get a quick feel for the interaction of these two laws by using the **Color** slider at the bottom of the screen.

Law of Closure. The Law of Closure is a great benefit to people who have had to try to decipher what I have written. A task which you mercifully have been spared. Open Interactive Illustration 5.x, Law of Closure [link to media]. In this illustration, we will get a chance to see how closure works, and some of the limits of this law. The first stimulus will be a box. Using the **Stimulus** drop down menu, you can select a square or a circle. You will have a **Gap Size** slider at the right side of your screen. I will describe how the slider works with the square, but it works the same way for all of the shapes.

When you first load a shape, the shape is closed and easy to be seen. Now use the **Gap Size** slider and add some gaps to the four edges. The separations will start from the middle of the side and as you increase the size, the gap edges move toward the corners. When the gap is at its largest, only the corners will be visible. As long as the gap is not too large, you will still see the pattern on the screen as a square. The small gaps, even though they disrupt the line, do not at all interfere with the perception of a completed square. You will see the gaps, so closure is not related to acuity. As you increase the gaps and the edges get smaller you may reach a point where you cannot see enough of the line to still perceive the box. Instead you will see disembodied line segments, but to get to this point, the gaps have to be quite large.

It may now be clear what role this grouping tendency in our perception has in the reading of handwriting. Often people will not completely close their o's or p's or 0's or other letters and numbers that

have loops in them. If we did not have the strong tendency to close the gaps in these letters and numbers, it might be a lot harder to read handwriting than it is. You might ask why we have such a strong tendency to close off figures. Well, consider all the times you look at an object that is partly blocked or occluded. Think of a table blocked by a book or computer or chair, or a house behind trees. Think of a predator hunting an animal in a forest, or a prey trying to avoid a predator in the forest. If our perceptual system was incapacitated by these breaks, we would have a really hard time making out objects in the real world (REFs and examples to build this section).

You can use this interactive figure to see how the law of closure interacts with some of the other laws that have been covered in this section. Use the **Color** slider to make the upper right and lower left quadrant of the figure a different color. How is closure affected by the differences in color? You can change the size of the figure, which will alter the proximity of the different parts of the figure. Is closure easier or harder with a large figure?

[possible interactive figure: pictures of scenes with bar occlusions – slider to reduce the amount of occlusion – what is necessary to make out the figure – give feedback in terms of the percentage of the scene - have randomly pic – load separately.

Possible pictures: building, face, forest, river, hall, field, quad

Could turn this into an experiment].

Law of Symmetry. Ever notice that regular and symmetric figures are easier to perceive, stand out, and are often said to be more pleasing? Perhaps that is due to the fact that our perceptual systems seem tuned to group objects based on the symmetry in the pattern. More symmetric edges are seen as belonging together. Open Interactive Illustration 5.x, Law of Symmetry [link to media]. In this illustration, there are dark and light vertical regions. The edges that border the darker areas are symmetrical about the vertical line running down the middle of that dark region. As a result, the dark regions appear as the figure, while the lighter region appears as the background. As you move the **Symmetry** slider on the right side of the screen to a lower position, the lighter edges will become more symmetrical, and the edges around the dark regions will become less symmetrical. As you do this, you will eventually reach a point where the light regions become the figure and the dark regions will become the background (Wertheimer, 1923/1938). This figure is made with a lot of randomness built into the drawing of the figure. You can

press the **New Pattern** button, and a new random set of shapes will be drawn, and you can see that the effects of symmetry depend really on the symmetry and not the characteristics of some cleverly developed set of figures.

Law of Common Fate. All of the laws that we have examined so far have dealt with static or still objects. There are some aspects of form perception, according to Gestalt Psychology, that arise only in the presence of motion. The Law of Common Fate describes one such situation. Open Interactive Illustration 5.x, Law of Common Fate [link to media]. In this illustration, we will examine how motion can by itself cause us to see objects as grouped together. On the screen will be a large number (60) cyan dots moving randomly. Among these dots, there will be a group of 5 moving together from the top left to the lower right on the screen. All of the dots are moving at about the same speed over all of the screen. The random dots are moving in random directions, as well as moving randomly relative to each other. The five dots moving in the same direction and together, however, appear grouped rigidly because all of their motion is the same as each others. The motion of the group of five does not need to really be any different than the random dots' motion. What is important is that they really move the same relative to each other. You can make this observation by using the **Direction** slider on the bottom right side of the screen. This slider controls how much the motion of the five dots is biased to move in this top/left to lower/right direction that you see when you first bring up the demonstration. By moving the slider down, the movement of the group becomes more random. Move the **Direction** slider all the way to the bottom and the motion of the group of five dots will be random just like the rest of the dots, but all five will move in exactly the same random pattern. Try it and see if the five dots are not still perceived as a group. The common movement among the five dots is key to this perceptual grouping. You can observe this by altering the amount of the movement among the five dots that is actually correlated with each other. You can adjust the strength of the correlation of the movement of the dots using the **Correlation** slider on the top right side of the screen. When the **Correlation** is at the top of its slide range, all of the dots move exactly the same. When at the bottom, the movements of the five dots are completely random relative to each other (but limited by distance, see below). When you make the movements of the five dots random

relative to each other, they are not seen as a separate group and hard to pick out from the back ground. You might make it random and ask a friend if they can find the five dots that form a group. You can select a new random set of dots using the **New Pattern** button at the top right of the screen. The dots are always randomly selected, so you can see that the grouping does not depend heavily on where the dots are relative to each other.

You can manipulate many other features of this illustration to see how the common fate law holds up and interacts with other laws. For example, you can make the five dots be a different color using the **Color** slider, again at the top right side of the screen. If the dots are a different color, how does this similarity aid in grouping the dots? Do the dots still need to be moving together to be seen as a group? You can adjust the range that the five dots can be from each other using the **Distance** slider. Even when the dots move randomly, they will never be farther apart than the value indicated by the **Distance** slider. As you move the slider up, the range of distance that will be allowed both for random motion and for new patterns will increase. You can only adjust this distance value when the correlation slider is not all the way at the top; since by definition, when the correlation is perfect this distance slider will not have any impact on the five dots. The **Density** slider will adjust the number of dots on the screen. Is the grouping easier to see when there a lot or a few other dots? Finally, you can adjust the speed of the motions of the dots using the **Speed** slider. Do you still see the group when the dots are moving very quickly across the screen? What do you get the impression of: the group, or can you still make out the individual dots? This distinction between the perception of motion and being able to resolve individual elements of an object will be important in the chapter on motion (Chapter 7).

Texture Segmentation

[I am having a hard time deciding what to do with this section.

<u>Good Figure or Prägnanz</u>

All of these laws, in some sense, fall under an overall Gestalt law: the Law of Prägnanz or roughly in English, the Law of Good Figure. According to this law, our perceptual system organized our perception to see things in the simplest and most regular fashion. Take the law of good continuation. To see two lines crossing is a simpler, more regular feature than to see four lines coming together, or two lines with sharp angles that meet. Similar arguments can be made for the other laws of perception developed by the Gestalt psychologists (Wertheimer, 1923/1938). Another aspect of the law of prägnanz is that the most stable way of perceiving a scene is to be preferred. In this case, they are referring back to what makes an object stable in the physical world. *[how can I give an example here?]* For example, in the law of common fate, objects that move together are rigid and more stable than an object composed of parts moving in random directs. *Ugh.*

[check with the Wertheimer Article].

Higher Level Processes

So far, the discussion of how form perception proceeds has assumed that everything about our perception of a stimulus will be determined solely by what is in the image. However, we saw several examples in Chapter 1, and even some in this chapter, that indicates that the forms we see are not so completely determined by the stimulus characteristics. Issues such as experiences, context, expectations, and even attitudes can alter how we see an object. In this section, we will examine some of these issues in form perception.

Top-Down versus Bottom-Up Perception

On of the most fundamental distinctions in perception is between Top-Down and Bottom-Up influences on perception. The discussions of contours, spatial frequency, and even most of the Gestalt psychology represent bottom-up perception. Here, the perception of the object was controlled by the nature of the stimulus. In other cases, our expectations, ideas, the surrounding stimuli, and other more cognitive factors can influence and even radically alter our perceptions. Open Interactive Illustration 5.x, What is this Figure? [link to media NOT DONE] You, hopefully, will remember this figure from Chapter 1. As you go through the demonstration, the figure will take on the characteristics of either a 13 or a B, depending on the context. It is important to note that the stimulus itself does not change in even the slightest degree. The change is in your interpretation of the stimulus. This example illustrates just one way

that your perception depends on more than just the stimulus. In this section of the chapter, some of the ways that our cognitive operations affect our perception will be explored.

Ambiguous Figure Ground Perception. Examine Figure 5.x. Is the picture a representation of a vase or a pair of faces examining each other? Actually, it is easy to see the image either way. This is an example of what the Gestalt psychological perspective would call an ambiguous figure-ground image. There are easily two different ways to see the image. And it is possible to change from perceiving one form to the other perception under conscious control. But it is not possible to perceive both forms of the image simultaneously. A few examples will help illustrate this phenomenon. The artist M.C. Escher, among other interests, enjoyed playing with the concepts of ambiguous figure ground relationships. In his woodcut, Sky and Water I (1938), the top and the bottom of the figure seem normal, with birds flying in the sky at the top and fish swimming in a sea at the bottom (Figure 5.x). However in the middle of the figure, the fish and birds meet. The fish form the background for the birds, and vice versa. Here you can either see the light fish or the dark birds, but not both at the same time. You have to switch back and forth to see both. It may take a while to see one of the animals. For example, when you start at the top of the figure and see the birds, when you look at the middle row, you might not notice the fish for a while. Let us explore this idea a little bit further and try a brief demonstration.

Open Interactive Illustration 5.x, Describe this Face [link to media NOT DONE]. In this illustration, you will be presented one image first, and then see how you would describe the second face that you see. When you open the screen, it is black. When you are ready to run the experiment, press the **Start** button. Run this illustration and then return to the book.

In this illustration, did the second image look like a young woman looking away from you or an old woman looking forward and down to the left? Actually, this second image can be seen either way. Try it for a minute and see if you can see both images. However, the first image was clearly one or the other face, and was probably biased which image you saw first in the ambiguous image. The program randomly picks which face you see first, so you can run the experiment again to see if the first image really does bias the first perception of the ambiguous image. However, this experiment does suggest that in images that can be interpreted in more than one way, how you see the image can be biased by your experience. This relates to the next topic, expectations.

Expectations or Perceptual Set. What we expect to see in a certain situation can influence what we actually see. In **Interactive Illustration 5.x**, **Describe this Face**, whether you tended to see the old woman or the young woman in the ambiguous figure, depended upon the expectancy developed by the first face that you viewed. There are many examples that can be shown that illustrate this effect. Let us try an experiment. Open **Experiment 5.x**, **Expectancy [link to media NOT DONE]**. In this experiment you will be briefly presented with a stimulus, and you are to type on the keyboard what the stimulus is. When you are done, press the **Done** button. There will be X trials. Do the experiment now and then return to the text.

Now that you are done, press the **Results** button. You will see two tables, each with two rows and two columns. One table is labeled with an "**I**" and the other table is labeled with a a "**I3**"[to me find a better font or symbol]. These represent the stimuli that were presented to you. Let us look at the "**I**" table first. The two columns are showing the letters that were immediately preceding the character. In one column, the characters preceding the target character is a capital letter, and in the other column the immediately preceding letters were lower case. The rows are labeled "**L**" and "**i**" Each cell represents the proportion of times you {or all students?} judged that "**I**" to be either an L or an I. Notice that when the preceding characters are upper case, you are much more likely to make the judgment that the character is an I and when the preceding characters are lower case you are much more likely to see the character as an L. You see a similar pattern for the "**I3**" character. In this case, the preceding stimuli are either capital letters of a similar font or other two digit numerals. Again, you tend to see the "**I3**" as a B when it follows capital letters, and as a 13 when it follows numerals.

The preceding characters have set up an expectancy. It is this expectancy that has influenced the way that you perceive these stimuli. The stimuli are always physically identical, but you interpret them differently. This tendency to see based upon your expectancy has been called a perceptual set (REF).

<u>Context.</u>

<u>Attentional Processes.</u>

a. Change Blindness

- b. Include some discussion of stabilized images call the section:c. Eye Movements and Form Perception



Figure 5.x. The Limulus or horseshoe crab with its compound eye shown (Grabbed from <u>http://www.mbl.edu/animals/Limulus/index.html</u> no permission yet).



Figure 5.x Vases or Faces?



Figure 5.x, Sky & Water I, 1938, Woodcut, M.C. Escher (permission not yet obtained).