# Chapter 4

# **Fundamental Limits of Vision**

# **Chapter Outline:**

- I. The absolute threshold of Vision
  - a. Overview of the Experiment
  - b. Determining the environment
  - c. Determining the stimulus
  - d. Determining the threshold.
  - The Duplex Theory of Vision
    - a. Spectral Sensitivity and the Purkinje Shift
    - b. Spatial Summation and Acuity
    - c. Temporal Summation
- III. Dark and Light Adaptation
  - a. The Dynamic Range of Vision
  - b. Scientific Notation
  - c. Adaptation in General
  - d. Dark Adaptation
  - e. Light Adaptation
  - f. Dark Adaptation and the Purkinjie Shift
- IV. Acuity

II.

- a. Hyperacuities
- V. Eye Movements
  - a. Types of Eye Movements
  - b. Version Eye Movements
    - i. Eye Muscles
    - ii. Saccadic Eye Movements
    - iii. Saccades and Heads Up Displays (Huds)
    - iv. Saccadic Suppression
    - v. Smooth Pursuit Eye Movements
  - c. Vergence Eye Movements

#### Chapter 4

#### **Fundamental Limits of Vision**

Our senses give us a great impression of immediacy. We see, hear, touch, smell, and taste with great speed and little apparent effort. This impression of immediacy leads us to operate on the assumption that our senses work automatically. However, as was indicated in the introduction, sensation and perception are complex activities. In this chapter, some of the basic features of vision will be covered to give some indication of the complexities of the act of seeing. Let us look at a couple of questions right up front to set the stage for the chapter. For example, what does it take to be able to see the dimmest light that we can see? Even this question reveals some unexpected complexities. The chapter will start with this question about how we see can be phrased: how big does something have to be for me to be able to see it? I have been purposely vague with this statement, because it turns out that the way that the question is asked will change the answer. The final topic in this chapter is a discussion of eye movements. That may seem like an odd final topic to be added here, but one very important reason we need eye movements will be apparent after we cover the question about that smallest objects we can see.

#### The Absolute Threshold of Vision

The first topic to be covered in this chapter is the absolute threshold for vision. The fundamental determination of the absolute threshold dates back to the middle of the last century (e.g., Hecht, Schlaer, & Pirenne, 1942). The way I want to approach this topic is to describe many of the choices a researcher needs to consider in order to measure the absolute threshold for vision. By examining these decisions, we will touch on many of the topics in this chapter, and it will become clear why these topics are important. My approach is modeled on a classic text in visual perception by Tom Cornsweet (1970).

#### **Overview of the Experiment**

The basic structure of an experiment for determining the absolute threshold has already been covered in Chapter 2 on psychophysics. Imagine a person being presented a brief stimulus using the Method of Constant Stimuli. Before this experiment can be run, the experimenter will need to know what the conditions of the room should be, e.g., how light or dark the surroundings should be. In addition, the researcher needs to know all about those stimuli that will be presented to the subject. Recall that in psychophysics, to get clear answers to our research questions, the onus is placed on tight control of the stimuli. So exactly what the characteristics of the stimuli ought to be, has to be carefully determined. **Determining the Environment** 

In measuring the absolute threshold, the researcher is interested in a very specific value, the ultimate limit of sensitivity of the visual system. That may seem obvious, but this value is not the dimmest light you could detect right now if you were suddenly placed in an absolute threshold experiment. To see how the absolute threshold is different from the dimmest you can see right now, consider walking into a dark building on a bright spring day. I am reasonably confident that your campus has some of these old dark buildings around. At first, the entire building seems dull and dark. Then, after awhile it seems normal. Now if you reverse the process and go outside, at first the outside may be painfully bright. Then after a brief period of time, the outside is no longer painful and appears normal as well. This simple example informs the researcher that our visual system adjusts or adapts to different levels of light. These are the processes of dark adaptation and light adaptation [both words are in the glossary]. Dark adaptation is the process in which the visual system increases its sensitivity, so that it can better detect low light levels. In the example above, dark adaptation had to occur to be able to see normally inside. Light adaptation is the reverse and happened when going outside in the example above. So the experimenter will want to have the person dark adapted completely, making the participant maximally sensitive to light. The participant must be put in a completely dark room for a period of time. It would be nice to know how long to put the person in the dark room. That will be discussed later in the section on dark and light adaptation.

To review, one feature of the environment needed for an absolute threshold experiment is a dark room, and the participant cannot just jump in the room and start the experiment. The person will have to sit there awhile.

A second feature of the environment is that the persons head will have to be in a fixed location. We need to know where the person is looking so that we can get the stimulus to flash into the eye. So the head should be held still, perhaps using a head and chin rest.

#### **Determining the Stimulus**

There are several characteristics of the stimulus that are relevant to doing an experiment to determine the absolute threshold.

The first characteristic of the stimulus to consider is where on the retina to present it. Recall from the last chapter that the rods and cones are not evenly distributed across the retina. Examine Figure 3.x **[The figure with the density of rods and cones on the retina – can we copy this here?]**. In some regions, such as the fovea about 20 degrees into the periphery, the receptors are most dense. Where the receptors are more dense, generally the eye is more sensitive to light. But the two types of receptors are not equally sensitive to light. Rods are much more sensitive to light. Since our subject is dark adapted, the subject will be using their rods to detect the light when the light is near our absolute threshold. This fact rules out presenting the stimulus to the fovea. A stimulus presented 20 degrees off to the side will hit the region of the retina where there are the most rods, which should give our participant the best chance to detect the stimulus. This difference in sensitivity between the rods and cones is part of what is called the **Duplex Theory of Vision [to glossary]**.

The next feature of the stimulus to consider is how long to make the flash. It may seem that seeing occurs instantaneously, but our retina actually collects light for brief periods of times to allow us to see. In some small ways, our retina is like film. One similarity is that the retina is like film which needs to be exposed for a brief period of time. To have a picture be bright enough to see, the shutter stays open for a brief period of time – the lower the light level in the scene, the longer the time period the shutter needs to stay open. So the time period that the eye collects light is analogous to the time period that the camera shutter stays open. Examine Figure 4.x. The picture on the left (a) was taken in daylight and the shutter is only open for  $1/30^{th}$  of a second. The picture on the right is a picture of Venus (the bright "star") and Mercury (the dimmer "star" nearby) taken when they were near each other in the morning sky. To get this image, the shutter had to be open for 1 second, and still the overall picture is darker than the daytime picture.

Now, the eye is not a camera and the retina is not film. The eye does not have a shutter and because we have dark adaptation, we don't need as much light in the dark to see as we do in the day. However, the time period that light needs to be collected by the eye is still not instantaneous. If we turn on the flash for too short a period of time, we will need more light to reach threshold than if we use a longer duration flash. This is known as **temporal summation [to glossary]** (Karn, 1936; Long, 1951). So, if we really want to know the dimmest flash that can be seen, it is important to make sure this flash is longer than the temporal summation period.

There is a similar phenomenon dealing with the size of the stimulus, and it informs the next characteristic of the stimulus to consider. Open **Experiment 4.x, Spatial Summation [link to media figure]**. In this experiment, the screen will be divided into two different regions. On the left half of the screen is the stimulus area with a red fixation mark in the middle of the screen. On the right half of the screen are the controls for this experiment, and a data graph that will plot your results as you make your judgments. Stare at the center of the fixation mark. There is a dot towards the top of your screen. It is really there – just very small and the same light level as the background. You can adjust the **Adjust Dot** 

**Brightness** slider on the control panel on the top of the right half of the screen and see it. What you are going to do is a JND type of experiment using the method of adjustment. The independent variable is the size of the dot that you will be adjusting. It starts out very small and then gets progressively larger with each trial. Each dot size makes up one trial.

Focus on the center of the fixation mark and drag the **Adjust Dot Brightness** slider until you can just barely see the dot. As you adjust, the graph will be plotting your judgment right along with you. When you think you are right at your threshold (you can just see it), press the button labeled **Threshold** right below the slider. Then the dot will go back to the background and go to the next size

which is larger. Repeat this task until you have done all of the trials, and then come back to the text. **[to production – can we collect this data from the students?]** 

Now that you have collected the data, let us see what we can make of it. If your data are similar to the usual pattern, for the smaller dots it takes a lot more of a light increase to see them than it does for the larger dots. But, it is not quite that simple, because as the dots get larger, the seems to be a point where the

increases in dot sizes no longer lead to decreases in the threshold for the dot. Just as we add up light for very brief periods of time, we add up the light over small regions of the world. If the dot is smaller than this region, then you need to have more light in the dot to make up for the small size and make the dot visible or even bright. When the light is larger than this region, making a light larger no longer makes the dot easier to see. This pattern of findings is called **spatial summation [to glossary]**.

So far, we know that the stimulus will be presented 20 degrees to the side, and for a duration longer than the interval of temporal and spatial summation. There is one final characteristic to consider here for our purposes. First open up **Experiment 4.x, Spectral Sensitivity [link to media figure]**. When the screen comes up, you will see five different colored circles, each with a slider to the right of the circle. This slider controls the intensity of the light in that circle. The slider for the middle, green circle is fixed and cannot be moved. Now adjust the other four so that they match the green circle in brightness. I suggest starting by changing the two circles next to the green (the yellow and cyan or pale blue) so that they match the green in brightness and then doing the outer circles (the red and blue). When you are done, return to the text.

Now, as you will notice, the sliders are at different positions for each of the five colors. Usually the slider for the green circle is the highest on the screen and the slider for the other colors are positioned lower, with red and blue having the lowest sliders. While this experiment is just a simulation of what is really happening, it does convey something meaningful about our visual system. We are not equally sensitive to all wavelengths. This difference in sensitivity across wavelengths is called **spectral sensitivity [to glossary]**. We are more sensitive to the wavelengths near the center of the visible spectrum than those wavelengths the end. Now, in this simulated experiment you used your cones, which you can tell because you saw color differences. However, because a participant is dark adapted in an absolute threshold experiment, the rods will be the receptors that will be functioning in the eye. Therefore, it is necessary to determine the spectral sensitivity of the rods for this experiment. When this is determined, it makes most sense to use the wavelength of light we are most sensitive to.

#### **Determining the Threshold**

Table 4.x summarizes all that we have been through so far. This table presents a review of the conditions we want for our experiment. When these conditions are met, and a few others besides, then the research will be in a position to determine the absolute threshold for vision. As would be expected, this value is going to be very small. If we only concern ourselves with the light that falls within a temporal summation period and spatial summation area, then the number of those photons of light that need to enter the eye to reach threshold are under 100 quanta (Barlow, & Mollon, 1982; Hecht, Schlaer, & Pirenne, 1942). To make this value for absolute threshold somewhat concrete, consider that there are in excess of 1 billion quanta being emitted from a tungsten light bulb in the same brief period as was used in an absolute threshold experiment. Obviously many other animals function much better in the dark than we do. We shall discuss this feature of their vision when dark and light adaptation are discussed.

#### The Duplex Theory of Vision

Hinting in the background of a number of the items that we have discussed is a strange and wonderful feature of our visual system. This is the idea that our visual system can operate in fundamentally different ways, depending upon the conditions in the environment. Recall from the last chapter that we have two different types of receptors, the rods and the cones. When there are differences in the structures of our body, as there are between the rods and the cones, it is expected that the two structures will operate differently in some way. This principle holds both for large scale structures like organs, and for small scale structures like cells of the nervous system. So the rods and cones seem to operate differently, and we have already come across or hinted at a few of the differences before. Below is a summary of those differences:

Rods are more sensitive to light overall than cones, which is why we were testing the rods for dark adaptation.

Rods are most sensitive to different wavelengths than cones. That is, they have a different spectral sensitivity.

Rods and cones have different spatial and summation properties.

Cones support color vision but rods do not.

With these differences, researchers have proposed what is called the **duplex theory of vision [to glossary]**. With the duplex theory of vision, it is argued that our visual system operates in two distinctly different ways, called **photopic vision [to glossary]** and **scotopic vision [to glossary]**. The photopic visual

system is associated with the operation of the cones and the scotopic visual system is associated with the operation of the rods. Table 4.x summarizes the differences between these two visual systems.

The table gives some of the basic features of photopic and scotopic. In general, we are much more familiar with photopic vision, as that is what is operating during the day time. As you can see from the table, during photopic vision we have color vision, we are most sensitive to a wavelength near 550 nm, we have good acuity (at least in the fovea) and we have relatively short temporal summation. During the night we are most sensitive to wavelengths that are nearer to 505 nm, but in general we have poorer vision with the scotopic visual system: less acuity, slower temporal summation, and no color vision. Open up **Animation 4.x, Simulation of Going from Photopic to Scotopic Vision [link to media]** to see a simulation of what these changes do to our visual abilities.

The first image that comes up is a picture of some children playing a soccer game. It is a color image. Now, the first factor that was mentioned as missing from scotopic vision is that there is no color vision. Click on the link below the photograph, and a short movie will come up that will show what the image looks like without color. You can replay the movie as often as you like using the movie controls right at the bottom of the movie image. The reason that scotopic vision lacks color vision is because it uses rods and there is only one type of rod. With photopic vision, there are three classes of cones. It is the fact that there is more than one type of cone that enables color vision. This matter will be discussed further in Chapter 6. A new browser window has been opened. Keep it open. Other activities will be interspersed but we will keep returning to this activity.

#### Spectral Sensitivity and the Purkinje Shift

Referring back to Table 4.x, the next item is the difference in the peak sensitivity of the rods and cones, or their different spectral sensitivity. Spectral sensitivity refers to the relative sensitivity of a receptor type to all of the wavelengths. There is more to the difference in spectral sensitivity between rods and cones than just the observation that cones are most sensitive to about 555 nm and rods are most sensitive to about 505 nm. Examine Figure 4.x. This is a graph of the spectral sensitivity functions for both the rods and the cones. This figure can be a bit misleading if you do not look at it carefully. Notice that the y-axis label reads Relative Sensitivity. Thus, each curve is plotted with the wavelength that they are more sensitive to at the level of 1 on the graph. By looking back at Table 4.x, you will see that the rods are far more sensitive to light than the cones overall. So if the absolute sensitivity of the two curves were plotted, then the cone curve would be much lower than the rod curve. This way of plotting the graph highlights the differences in spectral sensitivity between the two graphs. As can be easily seen, the two graphs peak at a different locations as indicated by the different wavelengths of peak sensitivity from Table 4.x. Notice also that the rods are relatively more sensitive to the short wavelengths, while the cones are relatively more sensitive to the longer wavelengths. This difference in spectral sensitivity is named the Purkinje Effect [to glossary] after the physicist {NAME} Purkinje who provided an early description of the phenomenon. These differences in the relative sensitivity of rods and cones to long and short wavelength lights is simulated in Interactive Illustration 4.x, The Purkinje Shift [link to media] (REF). In this illustration, you will match the intensity of two color discs. Since the discs are in color, they are using the photopic system. You will make the match using the slider on the top right hand portion of the screen. The two discs below are simulations of the relative brightness that the two patches would have in scotopic vision. They are both shades of gray, because the scotopic system does not see in color. The left hand patch is always from a relatively short wavelength color, and the right hand column is always from a relatively long wavelength color. After you have made a match, you can click on the **Match** button and see how different the two scotopic patches are. Pressing the **New Colors** button will randomly select two new colors to match for the top two discs. As you can see from trying this task, the right hand color will always be darker on the scotopic side when the two photopic colors match in brightness.

Going back to the Interactive Illustration 4.x, Simulation of Going from Photopic to Scotopic Vision, it means that merely making the photograph black and white will not adequately simulate the way that the scene will look to the scotopic system. The relative brightness of the regions associated with the different colors needs to be changed. Going back to the browser window with the photograph, click on the link that says **Simulate the Purkinje Effect**. As you watch the simulation, notice how the shorts and sock of the white team, which were red, become very dark and almost black, while the blue team's uniforms stay relatively bright. This simulates the Purkinje Effect.

While the topic of the Purkinje Effect is current, it seems a good time to mention the reasoning behind the move to eliminate of an American icon, the red fire truck. Many of your faculty and parents

grew up with that one color for the fire truck. Then in the 1970's [CHECK DATE], chartreuse fire engines started being made, much to the consternation of many Americans. However, from your reflections on the Purkinje Effect, you can see that a red fire engine might be very hard to see at night, and an emergency vehicle that is hard to see might lead to some safety hazards. Currently, there are many colors of fire engines, with lots still being the familiar red. The emergency lights are probably thought to more than compensate for the blackness a red vehicle would have at night.

### **Spatial Summation and Acuity**

The next difference between photopic and scotopic vision listed in Table 4.x is the relative acuity of the two systems. Acuity refers to the ability to see or resolve fine details, and it is inversely related to the size of spatial summation. As discussed above, spatial summation refers to the region where light is added together across a region of the retina. This addition effect is summarized by the equation:

K=I x A

This equation states that the brightness of a spot is determined both the intensity (I) of the flash and the size (A) of the flash. So, doubling the area and doubling the intensity of the dot will have the same effect on the brightness of the stimulus. However, this spatial summation only goes up to a limited area on the retina; for scotopic vision, the spatial summation region is about 10 arc minutes. In many cases, the size of a visual stimulus is given as the size of the angle the object forms on the retina. An arc minute is a unit of measure of the size of a stimulus on the retina. In every degree of angle there are 60 arc minutes, so the size of region where this summation takes place is not very large. But the region of spatial summation for photopic vision is only X (I think 1 goes here) arc min, which shows that in the daytime spatial summation is over a much smaller region of the retina. Within a region of spatial summation, the eye cannot make any discriminations, or at least not clear ones. So the range of spatial summation is related to our acuity. Thus, in the daytime, our acuity is going to be much greater than at nighttime. This difference in acuity needs to be accounted for in the photograph that is being used to simulate the change from photopic to scotopic vision. Click on the link in Interactive Illustration 4.x, Simulation of Going from Photopic to Scotopic Vision that says **Reduce the Acuity** below the picture. The picture will be blurred approximately to the same extent that it would be in the evening. Notice how much blurrier the image is after you have run the simulation.

#### **Temporal Summation**

The final feature that will be simulated is the temporal summation of the scotopic versus the photopic visual systems. Temporal summation refers to the period of time during which light is added up to make a single response. Thus, two flashes that occur within this time period will be experienced as a single flash. This feature of our visual system can be summed up, pardon the pun, in the following equation, knows as Bloch's Law:

 $K = I \times T$ 

(2)

Just as with spatial summation, all of the light in the brief period of time is added. This period of time is about 100 msec for the dark adapted rods (REF). This time is much shorter for the cones in the daylight and is about X msec. This temporal summation is a very important value to us and relates to another value talked about a lot, the **Critical Fusion Frequency (CFF)[to glossary].** This CFF is the frequency at which a flickering stimulus is seen as a continuous stimulus. If stimulus is flickering faster than this frequency, then the stimulus will be seen as a continuous stimulus. In essence, successive flashes are faster than the temporal summation period. For our photopic vision, this value is about 60 Hertz (Hz) or 60 flashes a second.

There are many flickering stimuli around us that take advantage of this feature of our visual systems. For example, movies and televisions are flickering stimuli and the rate of their flicker is carefully chosen by taking into account our CFF. This is an example of an area called Human Factors, which is the application of knowledge of our human capabilities in engineering situations. Here, the CFF is used to make a flickering stimulus seem continuous. The flickering of these lights leads to a very interesting experience that you might have seen in a movie or TV or under some street lights, the Wagon Wheel Effect, where a wheel appears to be going backwards. Open up Interactive Illustration 4.x, Flicker\_Motion [link to media]. When you open the figure, you will be looking at a wheel in the middle of the screen. The wheel will be obviously flickering, but that will help us in our simulation. At the top of the screen is a button that reads **Show Dot** and text that shows the current speed of the wheel in degrees per every new drawing of the wheel, and a slider to control the speed of the wheel. The wheel will always rotate clockwise; that is, to the right. Now, move the slider. If you click on the slider, you should be able

(1)

to use the arrow keys to control it, which gives a more precise control of the speed. Moving the slider will rotate the wheel the number of degrees indicated by the text at the top. As you move the slider to the right, the speed increases, and for awhile, so does your perception of the speed of the wheel. Somewhere between 11 and 13 deg per update, the motion gets really odd looking, and then for faster motions, the wheel starts looking like it is moving backwards, ever more slowly, until it stops at 22.5 deg per update. If you continue to increase the speed, you will see the pattern repeat two more times until you reach the maximum speed drawn. The stimulus is a flickering stimulus, and the eye cannot see the motion of the wheel between the flash. So the visual system needs to make an inference as to what happened during the dark interval between the flash. In this case, since all the spokes of the wheel look identical, then the eye matches which spoke is closest when the wheel is seen again. The spokes are 22.5 degrees apart from each other, so when the wheel moves 22.5 degrees or a multiple of that movement, then the wheel looks still. Motions less than that value, but near to the multiple make the wheel look like it is moving backwards. It might be thought that this program is a trick, so try a variation of this demonstration. Press the **Show** 

**Dot** button. When this button is pressed, along one of the spokes of the wheel a red line will be drawn, you will still see the white spoke showing through. At the end of the red line a red dot will be drawn. This dot will always be drawn at the end of the same spoke. Now adjust the speed of the wheel to about 20 degrees per update. You will see the wheel going to the left and the dot and its spoke going around to the right. Because the dot and line are unique, they get matched up so that the motion will be perceived correctly, but the wheel will still be seen to be moving in the opposite direction. At 22.5 degrees per update the wheel will be still and the dot flying around. At faster speeds, about 25 degrees per update, you will see the real speed in the red dot, but the wheel will look like it is going much more slowly, say about 2.5 degrees per update. Can you figure out why?

Now, the temporal summation is over a lot longer period for scotopic vision than for photopic vision. You might have experienced the difference doing sparklers on the Fourth of July. In the daylight, if you move the sparkler around, you only see the sparkler. But at night, you will see a blur trailing after the sparker. This trail is due to the longer temporal summation of the scotopic vision. Open up Interactive Illustration 4.x, Temporal Summation [link to media] for a simple simulation of this phenomenon that might trigger your memory. When the interactive illustration comes up you can drag a blue dot around the window by clicking on the region white portion of the screen and dragging you mouse. The blue dot on the white background follows quite faithfully. Press the button at the top of the window and you will be in a scotopic mode. Now when you drag the dot, which is gray because we do not have color vision at night, the dot is spread across the screen, catching up if we hold the mouse still, just as the sparklers would. [To reviewers – should I cut this illustration – it strikes me as a bit cute – I could use some outside opinion here].

Now let us finish taking the picture of the soccer players to what it might appear like in scotopic vision. Return to the last image of the soccer players in **Interactive Illustration 4.x**, **Simulation of Going from Photopic to Scotopic Vision** which showed the poor acuity of the image. Click on the like below where it says **Simulate the slow temporal summation**. This simulation smears all of the images across the screen, assuming that the scene is moving.

Thus, scotopic vision is very different than photopic vision. In many ways it seems worse and if we had to rely on it during the day time, it would indeed be worse. We could not read, or keep up with events around us or even experience the colors we love so well. But remember, scotopic vision is for the nighttime. The larger spatial and temporal summation abilities are important because they help give scotopic vision some of the sensitivity it needs to operate then. If we only had our photopic vision, we would be in essence night blind. It is important to remember the functions of each type of visual system that we have.

#### **Dark and Light Adaptation**

Running in the background of a lot of the discussion so far is the notion that our visual system adjusts or adapts to the current lighting level. When measuring the absolute threshold, it was necessary to dark adapt the subject. This implied that visual sensitivity changes as a result of being in the darkness. This section will explore first the degree to which our visual sensitivity changes, and then the nature of the changes to our sensitivity will be discussed.

#### **The Dynamic Range of Vision**

Consider an object that weights about 1 pound. Something that weights a pound is pretty light. Assume that 1 pound value represents our absolute visual threshold. The heaviest weight that the average person can carry is about X pounds. Continuing the analogy between weight and light, the equivalent heaviest light that can be seen without causing damage to our eyes is approximately 1,000,000,000,000 or 1 trillion pounds. That is about 500,000,000 or 5 hundred million tons. To give you some perspective, the Battleship Missouri from WWII and the Korean War weighed only 58,000 tons fully loaded. The Empire State Building is 365 million tons.

### Scientific Notation [OUESTION: DO I NEED THIS SECTION]

Now these numbers are very large, and if they had not been converted to tons would have been unwieldy to handle. In the measures of light, there are not units like tons that the units of intensity can be converted to. In fact, the number ranges, like astronomical numbers, are simply too big. There is a simple notation that can compress these ranges of number to a more manageable size. This is called scientific notation and is a close analog to the use of logarithms that was discussed above. In scientific notation, the number is expressed as a value time 10 raised to some exponent as such:

 $1 \times 10^{12}$ 

If the value is a one it can easily be dropped. The range light intensities that can be seen by the visual system converted into scientific notation is, if a value of 1 is used for absolute threshold, from:  $1:10^{12}$ 

The 1 represents the absolute threshold, and the  $10^{12}$  represents the maximum safe exposure to light that the eye can handle. To translate the weight of the Missouri and the Empire State Building in these units, they first needed to be converted into pounds from tons and then converted to scientific notation. These are shown in Table 4.x. Notice how much more manageable these number are. In fact, if you use the logarithms, you will find that you end up with the same result.

I HAVE AN ACTIVITY TO ILLUSTRATE SENSE OF SCALE. IF I EVER GET IT RIGHT I MIGHT INCLUDE IT HERE]

#### Adaptation in General

Adaptation takes time. We do not instantly adjust to a new lighting level. The more the change in lighting level, the more time adaptation will take. You have probably had some experience with this phenomenon. Open Interactive Illustration 4.x, Dark/Light Adaptation [link to media] and it should help you understand what is being discussed. You may need to give this interactive illustration a little time to download. There is a picture that needs to be loaded, and depending on the speed of your machine and the network traffic, it might take a few moments. Depending on the java plugin that you are using, you may see a progress bar to tell you how the picture is coming along.

When the picture is loaded you will see the picture covering most of the screen and the picture appears normal. On the right hand side of the screen are three buttons: **Too Light**, **Too Dark**, and Adapt. There is also a slider labeled Light Level. The Too Light and Too Dark buttons, and **Light Level** slider, refer to the light level apparent in the picture relative to the current light level you are adapted to. Press the **Too Dark** button. The screen goes nearly black. How black will depend upon your monitor to some extent. You might have an impression of where the doors are but not much more. In this case, the simulation represents your coming in from a very, very bright outside and the inside is much darker. Now press the Adapt button. First, you will make out the door. Second, the lights on the second floor become visible as white dots with lines through them from the bars of the railing. As adaptation proceeds, more and more becomes visible, with the regions that are quite dark becoming visible last. This is a simulation of the experience you might have with dark adaptation, because you are adjusting or adapting your visual system to a less intense or darker environment. It is much faster than in the real visual system, but the gradually seeing more and more detail is quite realistic.

Light adaptation is the opposite procedure. You come into a very bright area from a relatively dark area, or you turn on the lights. You can simulate that on the Interactive Illustration by pressing the **Too Light** button. Now all is nearly white and you can only make out a few of the darker portions of the image. Press the **Adapt** button and you can see a simulation of light adaptation. Now it is the darker regions that are visible first and the brighter regions of the scene that only become visible later.

You can see how the amount of adaptation is related to the speed of simulation by using the Light Level slider. You can adjust the amount of adaptation to any desired level and press the Adapt

button and see how long it takes to adapt. Now you will notice that in general, light adaptation is faster than dark adaptation, a fact that will be discussed below, but other than that feature of adaptation; the more the adaptation, the longer the adaptation.

#### **Dark Adaptation**

Dark adaptation will be discussed separately from light adaptation, as it turns out that dark and light adaptation work rather differently from each other. We will start with dark adaptation since it was the first studied.

In a dark adaptation experiment, a subject will be first placed into a very bright room. To be able to measure how dark adaptation proceeds, it is necessary to light adapt the person first. Then, the person is plunged into complete darkness, and their absolute threshold is measured at repeated intervals over time, using one of the psychophysical methods discussed in chapter 2. Doing this type of experiment gives the pattern of data shown in <u>Interactive Figure 4.x</u>, <u>Dark Adaptation Curve</u>. The participant's absolute threshold falls rapidly when they are first in the dark. This is seen by the first part of the curve going downwards so steeply. Then their absolute threshold levels off. At about 7-8 minutes, their absolute threshold starts falling again, finally leveling off at the participant's lowest absolute threshold after about 30 minutes in the dark (Graham, 1965).

One of the interesting features of this graph is the leveling off of the threshold, and then its starting to fall again after about 7-8 minutes. When data show this kind of break, where it is not smooth like the rest of the graph, scientists begin to suspect that more than one type of process is involved. From our discussion of scotopic and photopic vision, we can make some good hypotheses about what could account for this pattern in the data. In the bright daylight cones are used, and in the dark rods are used. So the first part of the data should be the result of the cones and the second part of the data should be the result of the rods. Check the **Show Theoretical** check box on the top right of the screen to have this expectation shown. The cone function will be shown in green and the rod function in red. In bright light, rods are not as sensitive as cones so they do not work, and in the dark the cones are not as sensitive so they do not work. The question is how to test this prediction. The cone part is rather easy. The fovea, if you recall, only has cones, so stimulating this region of the retina only will allow us to see the effects of dark adaptation on just cones. You can see what the data look like in this type of experiment by clicking the **Test in Fovea** check box on the right hand side of the screen. The rod data disappears, and only the cone curve stays as would be expected. Testing the rods is a little harder, as there are cones throughout all of the retina. However, in the far regions of the periphery, there are a lot of rods and only a few cones. Press the **Test** in **Periphery** checkbox to see data from this region of the retina. The outcome will be more apparent if you clear the **Show Theoretical** checkbox by clicking on it again. Now the cone data is barely visible, just as would be expected. To really nail this prediction down, dark adaptation tests have also been run on rare individuals who do not have cones in their eyes. They are rod monochromats. You can see their data by clicking on the **Rod Monochromat** check box. Note the cone curve is completely absent.

Thus, it is clear that cones are responsible for the first portion of the curve and the rods for the second. However, these data do not explain how dark adaptation occurs, especially within the scotopic and photopic systems themselves. There appear to be several mechanisms that all play a role, but one of the more interesting facets is that it appears that the receptive fields themselves change in dark adaptation. It seems that the inhibitory surrounds of the center-surround receptive fields, at least in the retina, are much smaller and weaker when we are dark adapted than in the daytime (Lankheet, Row, van Wezel, van de Grind, 1996; Wist, 1976). Wist (1976) used an interesting procedure to test that the weakening of the surround antagonism of center-surround cells was involved. They used a stimulus called the Hermann Grid (Figure 4.x). Examine the figure and notices that the intersections, particularly those in your periphery, appear darker than the lines between the squares. This visual illusion is thought to be caused by the center-surround receptive fields, as will be discussed in Chapter 5. What Wist found was that the intersections of the Hermann Grid were not nearly as dark in dark adapted participants. These findings were very consistent with the lessening of the surround fields during dark adapted participants are findings in the surrounds of receptive fields is one of the reasons that scotopic vision has a larger spatial summation area than photopic vision.

# [should I find a way to illustrate the change in the receptive fields and what that does to an image?]

#### **Light Adaptation**

In many ways light adaptation is like dark adaptation, but in reverse. However, there is one major difference: light adaptation will be completed in about 5 minutes or even less. Dark adaptation is, therefore, much slower. It seems that this difference is partly due to the fact that light adaptation is driven actively by the light entering the eye, while dark adaptation is a more passive response to the lack of light. Also, in most cases dark adaptation is sufficiently fast. The changes from afternoon sun, through dusk, to evening is about the time it takes dark adaptation to complete its total swing.

#### Dark Adaptation and the Purkinje Shift

Yes, the rods are, overall, more sensitive to light than the cones. But recall the Pukinje shift. Look again at the spectral sensitivity function in Figure 4.x. The cones are relatively more sensitive to long wavelength light than the rods. The greater sensitivity of cone than rods to the very long wavelengths is so great that cones are actually more sensitive to these wavelengths than rods are in an absolute sense, not just the relative sense (Cornsweet, 1970). This observation has some interesting applications. If a task requires a person to go quickly from a lighted area to a dark area in a short period of time, then the period of dark adaptation could be a problem. This dark adaptation period delays how fast the person can begin functioning in the dark. However, this limitation can be overcome by having the person work in an environment with only red lights. The red light will not be absorbed by the rods very well, so they will begin to dark adapt. Then when the light is removed, the person is dark adapted to a large degree and can operate immediately in the dark.

A common example of this technique is the use of red flashlights when observing with a telescope at night. The observer can look at a star chart or other information using the red flashlight. This will give them the ability to use their photopic vision to read the details of the star chart. Then, when the light is turned off they still have their dark adapted rod vision to see the subtle detail in the night sky, such as a faint galaxy or nebula.

Another application of the knowledge of dark adaptation takes place in the modern aircraft cockpit. There are a lot of visual situations that can alter a person's ability to read a cockpit display. Think of driving a car with an electronic speedometer. In some cases, the numbers for the speed of your car may be very difficult to read. In the air, the situation can be worse. Think of a pilot flying into the sun. The pilot may light adapt to the sun and then not have sufficient sensitivity to read the electronic display when needed. I think we might all feel a little uneasy if we thought that pilots had to take a long time just to read their flight instruments. The opposite problem occurs at night, when it is desirable to not have the displays so bright that the pilots will light adapt to the display and not be able to see the runway clearly. To avoid these problems, researchers determined both the light level needed to see outside and the displays in these situations. From this data, they have developed automatic adjustments so that pilots don't have to spend a lot of time adjusting the luminance of their displays when the lighting conditions change, which might be because of flying into clouds or night (Krantz, Silverstein & Yeh, 1992; Silverstein & Merrifield, 1985).

#### Acuity

Our visual acuity refers to the ability to resolve fine details; that is, to see small features of what we are examining. As was mentioned above, our ability to see smaller and smaller objects is limited, among other things, by the size of our spatial summation region. It is also limited, in bright daylight, by how good an image our lens is capable of producing. Acuity is a crucial limit as it relates to our ability to read and pick up information from all of the images around us. In addition, there are many applications of this knowledge around us.

First, let us discuss the many types of stimuli that have been used to measure acuity. In all cases, there is some small scale feature of the stimulus that is relevant. It is the size of this feature when it can not longer be made out that is used to determine a person's acuity. The most popular type of acuity stimulus is the familiar Snellen Letters. These are the letters that you may have read in an optometrist's or ophthalmologist's office to determine your visual abilities. Variations of these Snellen Letters are quite common, and you probably read a variation of these letters when you took your driver's test. These letters are carefully designed and selected. The entire alphabet is not used, but only those that have needed features that can be confused. When that feature of the letter falls below your acuity limit, then you will tend to not recognize the letter or to confuse it with another letter.

Snellen stimuli are very easy to use in a screening situation, but they do have some limitations. One is that they require the participant to be able to read and probably speak English. Moreover, they do not give as precise a measure as might be wanted in some circumstances. As a result, several other stimuli have been developed. Open Interactive Illustration 4.x: Acuity and Stimuli [link to media] to see a sample of some of the more common stimuli used, and the critical feature that is measured to determine acuity. The stimulus on the left is a very common stimulus, and is called a grating. When the width of the bars falls below your acuity limit, you will see a gray field instead of the bars. So the critical feature is the bar width, indicated by the red bracket at the top of the grating. The checkerboard on the right side of the screen is very similar. Again, when the squares are too small, a grey field is seen. The critical feature then is the size of one of the squares. Both of these stimuli are said to measure resolution acuity. The middle stimulus is the Landolt C which is very useful for measuring acuity in people that cannot read. The critical feature is the gap in the circle. The way the test works is that the gap is either up, down, left, or right. The participant's task is to indicate the direction of the gap. The Landolt C and the Snellen acuity are said to be measures of recognition acuity (Riggs, 1965).

At the bottom of the figure, you can move the slider to adjust the size of the critical features of the stimuli. You can adjust them and see if you can reach your acuity limit for any of the stimuli. It is possible that you will not reach your acuity on these figures. The standard sizes of dots on these monitors are just below acuity thresholds at normal viewing distances. The way around this problem is to first adjust the critical features to their smallest size. Then move away from the screen and determine how far you need to be to make the features fall below your acuity. In fact, let me suggest a little experiment. After this first trial, make the features twice as big as they were and see now how far away from the screen you must be to be unable to see the critical features. Repeat this procedure a few times and see what happens. What do you think that implies about our visual ability?

Now, let us talk about measures of acuity. The most commonly used measure of acuity is in terms of **visual angle [to glossary]**. The visual angle is the size of an object at the eye. The goal of the measure is to find a value that will tell us how big objects will be at the eye when they can be resolved. The problem is that our normal measures for the sizes of objects will not work for the eye, because the size of an object depends upon the distance an object is from the eye.

Open up Interactive Illustration 4.x, Visual Angle [link to media] and we will explore the concept for a bit. When you open the interactivity, you will see the schematic eye to the left side of the screen with an arrow in about the middle of the screen. The arrow is red and with two lines that leave it, one at the top and the other at the bottom of the arrow. These lines meet at the front of the eye and then cross and form an angle that indicates the size of the object on the retina. There are two sliders that you can use to adjust the arrow. The slider at the bottom of the screen labeled **Distance 1** moves the arrow closer and farther from the eye. The slider at the right side of the screen labeled **Size 1** will change the size of the arrow. First, let us examine what happens when we change the distance of the object from the eye. As you move the distance slider, the physical size of the object does not change, but the angle formed by the arrow at the eye changes a lot. When the arrow is closer to the eye, the angle of the arrow is much larger than when it is at the opposite side of the screen, far from the eye. The angle is related to the size of that arrow on the retina, so when the arrow is close it causes a much larger image on the retina than when it is far.

To make the point even more concrete, click on the check box that says **Add Second Arrow**. When you do, the **Distance 2** slider and the **Size 2** slider will be enabled, as will two buttons **Match Physical Sizes** and **Match Angular Sizes**. When you press the **Match Physical Sizes** button it will adjust the two arrows so that they are the same length. Move one of the arrows close to the eye and the other arrow to the far side of the screen. It does not matter which one. Then adjust the far arrow to some size and press the **Match Physical Sizes** button. Observe how much bigger an angle is formed at the eye, and consequently on the retina, of the nearer object. Now press the **Match** 

**Angular Sizes**, and the closer arrow will be resized to cast the same size angle on the retina. Notice how much smaller the closer arrow is. Now, it is this angle that is the size of each object on the eye. Recall that our degree of spatial summation on the retina is related to our acuity and spatial summation is in terms of area on the retina. So for us to be able to see an object it is the area on the retina that is important to know. Our ability to see an object depends both on its size and its distance, which together determine the angle of the object at the eye. This angle is the visual angle mentioned above.

The normal adult is able to resolve an object or part of an object down to about 1 arcminute (1/60<sup>th</sup> of a degree). This value represents the normal acuity that is used when you take your driver's test or go to the optometrist. You may be familiar with the other measure of normal acuity, the Snellen Acuity, with a normal value in the United States of 20/20. This measure fixes the distance of observation so that the sizes can be known. Here is how it works. The top number refers to how far you are away from the object in feet, and it is always 20. The bottom number is how far away, again in feet, a person with normal acuity is able to resolve the same feature. So if the acuity is 20/20, both of you can make out the same objects. However, if you are 20/40, you need to be 20 feet away from an object to see what a normal person can see 40 feet away. This implies your vision is twice as bad as the normal observer. Try it on the figure. Move the red arrow to the most distant point you can and the blue arrow to about half way from the eye. Make the red arrow a reasonable size, and let us assume that this distance is 40 feet and this arrow is the smallest figure a normal person can see. Press the **Match Physical Size** button because in Snellen acuity both observers can see objects of the same size. Notice how much larger the angle of the blue arrow is than the red arrow. Now 20/40 is not very bad acuity. Imagine the situation for, say, 20/100 where the normal can make out a figure from five times the distance you have to be to see the same object.

### Acuity and Retinal Location

Go back to Figure 3.x which shows the distribution of cones across the retina. [Production: can we repeat the figure here] Since we are talking about daytime vision; we are talking about cones, and most of the cones are in or near the fovea. It might be likely that this pattern of cones would have some impact on our acuity. It seems to make sense that where there are more cones in the eye, you might have better acuity. Try Experiment 4.x, Acuity and Retinal Location [link to media]. In this experiment, you measure your acuity at five locations across the screen, while you are fixated on one edge. As we just saw, the visual angle that an object has depends upon the distance you are from that object. The same is true for where an object falls on the screen, for the very same reason; just consider where the points of the tops of the arrows fall in eye in the Interactive Illustration 4.x, Visual Angle. So for this experiment to be reasonably comparable to what others taking this test, we need to regulate how far you are from the screen. However, different screens are different sizes, and the stimulus will be in different places relative to the size of the screen. We can compensate for these differences by all being the same distance from the screen relative to the width of the screen of x times the width of your screen. Measure the actual size of the drawn area to get this value.

Once you are in place, hold your head still as best you can. You can check to make sure your head position is in the correct location periodically throughout the experiment. Now, open the experiment. This experiment has a new page that will precede most experiments in the rest of this text. On this first page, you will be able to adjust the features of the stimulus that you are going to use in the experiment. In this case, your options are very simple. You can choose the type of acuity target that you would like to use. There are three of the common acuity stimuli discussed above available for this test, a grating, a Landolt C, and a checkerboard. With the grating stimulus, you try to resolve the presence of the bars, as opposed to a gray region made by the spatial summation of the white with the gray bars. The critical feature is the width of the bars when you can just detect the grating. In the Landolt C, developed by XXXX (get REF), the task is to determine the direction of a gap in a circle, thus termed a C. You will indicate whether the opening is up, down, left or right. For the checkerboard, you are to try to resolve the squares on the board.

The second stimulus feature that you can manipulate in this experiment is the number of positions to test. During the trials, you will be looking to one side of the screen for fixation, and the number of positions refers to how many different locations, evenly split up, you will be testing across the width of the screen. The default is 5 and means that you will be determining your acuity 5 times, once each for 5 different retinal locations. The last one is the largest stimulus feature to test, the width of the bar for the grating, the gap in the Landolt C or the size of the squares in the checkerboard. The smallest feature will be the smallest that can be drawn by your monitor. Use the default values of the experiment. [which stimulus type should be the default?]

The next screen is the Method of Limits screen that you have seen before. Notice that the defaults here are changed. The number of levels to test is 9, and the default type of Method of Limits is the 1 up, 1 down Staircase. Again, use these defaults. The greater number of levels to test is necessary to get differences between some of the acuity positions.

button at the bottom of the screen.

To best run the grating or the checkerboard stimuli, it helps to perform a little calibration procedure to match the background to the brightness of what you will see when the eye cannot resolve the bars. The way a computer controls a monitor is not easily predictable from the values that are used to tell a computer monitor what to do. In addition, each monitor differs from every other monitor to a slight degree. So before running an experiment, it is often necessary to calibrate the screen to the individual user. When you can no longer resolve the bars of the grating or the squares of the checkerboard, it helps if the background is at a similar luminance level. What you will be doing in this case is adjusting the background so that it appears about the same brightness as a blurred grating. So, when you cannot resolve the grating or the checkerboard, it will disappear or nearly disappear. On this calibration screen you will have a slider on the bottom of the screen. Notice the grating at the extreme right side and look to the extreme left side of the screen. The grating has very narrow bars, and when you look to the left you will not be able to make out the bars. While looking to the left, use the slider to adjust the background so that it is similar to the brightness of the grating, making it as exact as you can. When you have your match, press the **Done** 

The next screen is the experiment screen. I will first give the instructions for the checkerboard stimulus, which is the default. The instructions for the grating are identical, so if you chose that stimulus you can follows these same directions. The instructions for the Landolt C are slightly different and they will be covered next. On the far left of the screen you will see the fixation mark. It is very important that you keep your eyes fixated on this location and do not move them. You will want to, so keep moving your eyes back to the fixation mark. In one position, the nearest, is a grating of the largest size. If you can make out the bars of the grating, press the **Yes** button or the **z** key. If you cannot make out the bars of the grating, press off of the fixation make to make your response. You can hold your fingers on the two keys.

You are doing a Staircase method starting with a descending staircase, so you should have no trouble seeing the bars of the grating. The next trial, the bars will be smaller until you can no longer see the individual bars of the grating, but they look like a gray smudge, more or less. Then the bars will get larger again until you can see them again. When the threshold for this first location has been determined, the procedure will start over but with a location now farther in your periphery. There will be five locations, each farther to the right of the fixation mark, so farther into the periphery.

The other stimulus that you might use to test acuity in this experiment is the Landolt C. Your job will be to indicate which direction is the opening in the circle. You can indicate the open direction of the circle in two ways: either press the arrows at the bottom of the screen, one for each of the four directions used, or press the arrow keys. If you don't know the direction of the opening, guess. Otherwise the procedure is identical to that for the grating stimulus.

When you have completed all of the locations, the acuity, a type of threshold, will be determine for each location and plotted on the screen in front of you. What you need to examine is, how does acuity change as the stimulus gets farther and farther into the periphery? Perform the study and then return here.

If your results are like most people's, and assuming you were careful in your observations and not too lucky in your guessing, you will see a linear increase in your acuity with the farther the stimulus gets into the periphery (Anstis, 1998). You can see this in **Interactive Illustration 4.x: Picturing Peripheral Acuity [link to media NOT DONE]**. This is a demonstration that is an interactive version of one developed by Anstis (1998). First, you need to be X times the width of the picture away from the screen. What you need to do, is examine the photograph in the center and then adjust the slider on the right until you just notice that the picture becomes blurred in the periphery. Then move your eyes and notice how much more blurred the picture is than you expected. Here is another way to try this illustration: move the slider to the top and make the periphery obviously blurry – then adjust it until you can no longer, but just no longer, see the blur in the periphery. The foveal region never becomes blurry, so in this picture no part of the image appears blurry – then look around. You will see the picture as being quite blurry. In some sense, this is how blurry our periphery is at all times. Only the foveal region is clear.

#### **Hyperacuities**

Open up **Interactive Illustration 4.x: Vernier Acuity [link to media]**, which will illustrate a different type of visual sensitivity: this time to location or at least change in location. When you open this

interactive illustration you will see two lines and your task will be to align them so that they are end to end and look like one line using the slider at the bottom of the screen. At the right end of the slider, how far apart from vertical the two lines are will be indicated. This offset value will be in pixels or the dots on the screen. You should be able to see the lines are actually lined up to within a small fraction of a pixel.

This task illustrated the exquisite sensitivity of your **vernier acuity [to glossary]**, which is your ability to detect when the two ends of lines are aligned. Your vernier acuity exceeds standard acuity at all retinal locations (REF ???). Try **Experiment 4.x**, **Measuring Vernier Acuity [link to media]**. This experiment is a simplified version of the last experiment that you ran testing acuity. In this case, your acuity object will be to detect when the two ends of a line are aligned. Again you will use the Staircase method and the measure will be in the same screen units as before. The question is, how does this vernier acuity compare to the acuity for the Landolt C object used above?

This experiment will proceed very similarly to the task above, except that only the first location will be used. You can select up to 5 positions, which will be the same as the default locations used in the last experiment. At the beginning of the experiment you will see the same type of stimulus setup window you saw at the beginning of the acuity and retinal location experiment. The main difference will be that the **Vernier Line** will be the stimulus default. You can also select the grating and Landolt C Stimuli if you wish to repeat these conditions from the earlier experiment.

Then you will go to the method window, which will be Method of Limits just like before. In this experiment, you are to indicate whether the lines look lined up or offset. The question is phrased "**Do** you see the break in the line?" You can press the **Yes** button on the window to indicate that you saw the break or the **No** button to indicate that you did not see the break. If you wish to use keys, the **z** key means **Yes** and the *I* key means **No** just as in the last acuity experiment. The data at the end of the experiment will show your vernier acuity in a bar chart. [Production: can we have the data from the same condition in the last experiment here as well?]

This experiment used the periphery to match the earlier acuity experiment. However, in the fovea, we can detect displacements in a line of about 7-8 arcseconds (an arcsecond is  $1/60^{\text{th}}$  of an arcminute as an arcminute is  $1/60^{\text{th}}$  of a degree of visual angle)(Klein & Levi, 1985; Westheimer, 1979). This is a far smaller acuity measure than standard values, which are on the order of 1 arcminute. This small value places vernier acuity in the class of acuities known as **hyperacuities [to glossary]**. We will run into another hyperacuity in Chapter 8 on depth perception.

The hyperacuities are very interesting for a couple of interrelated reasons: 1) they seem to suggest the ability to resolve differences finer than the region of spatial summation, 2) they suggest the ability to resolve differences that are even smaller than the width of a single cone, which are about 30 arcseconds in the fovea (REF). A break in a line about 1/5<sup>th</sup> the size of a foveal cone can be detected in the cones. A remarkable ability. However, it does not really contradict all that has been said about spatial summation, or even the ways that cones work. The two parts of lines are not in the same spatial summation regions, and the lines are larger than one cone. What vernier acuity implies is that there is an ability in the brain to integrate the signals across the activity of several cones and summation regions and see where the center or edge of a region falls with great precision. Our ability to determine location, which is what vernier acuity is, is greater than our ability to resolve detail, which is what acuity is.

You have all experienced the impact of your vernier acuity. Open Interactive Illustration 4.x, Vernier Acuity and Lines [link to media]. In this activity you will demonstrate your exquisite sensitivity to changes in location. The slider on the side of the screen simply rotates the line. You can rotate the line. When the line is not vertical, you can easily notice the steps in the line. When it is vertical, the steps in the line go away. See if you can make the line vertical. Simply rotate the line until you don't see any steps, and then press the **Vertical** button and see how well you have done. You will find the task very easy as long as you look up and down the line to notice the presences or absence of these steps. Press the **Reset** button at the bottom of the screen to get a new random starting position to try the task more than once.

These steps in the graphics community are given the technical term jaggies, really! The individual dots, called pixels on your monitor, are just below our standard acuity at a standard viewing distance. It is this reason that we do not see the individual pixels making up the line when it is vertical. However, the steps in the lines are well above our vernier acuity. A lot of work in the human factors area has been done to find ways to eliminate them from graphics, especially to find ways to take advantage of our visual systems to make it as easy as possible. The way that has been developed has been to antialias the lines.

Nontechnically, this means to blur the edges of the lines (Krantz, 2000; Silverstein, Krantz, Gomer, Yeh, & Monty, 1990). It was this knowledge that allowed the drawing of lines positioned to a fraction of a pixel in the first vernier acuity interactive illustration. Press the **Smoothed Line** check box at the bottom of the screen to smooth the line and you will see the same type of line as was used in that earlier demonstration. It will automatically pick a new randomly selected tilt for the line. Then, as for the unsmoothed line, see if you can still determine when the line is vertical with the same degree of accuracy. You will find that you are not quite as accurate as you were, though you will still be quite good.

#### **Eye Movements**

We move our eyes. A lot. A lot more than you even know. This is really a rather interesting observation. When you say to a friend, "Look at me!" you do not mean to be just in their field of view. In fact, when we say this line to someone, we often are in their field of view. You want to fall on a particular part of the retina. Based on our discussion of acuity, we begin to understand the need for eye movements: the limited range of the acuity. Really good acuity is restricted to the fovea. Acuity falls off rapidly even just outside of the fovea. The fovea is only about 1 degree in area. The full moon is about 30 arcminutes, so the fovea is approximately twice as wide as the moon. Thus, to see clearly, we must move our eyes so that whatever image we want to examine falls on the fovea. The periphery is used for other functions. We will run into some of these when we discuss orientation in Chapter 14.

#### **Types of Eye Movements**

Eye movements can be organized into many different types. There are voluntary and involuntary eye movements, large scale and miniature eye movements, and there are **version [to glossary]** and **vergence [to glossary]** eye movements. Voluntary eye movements are rather self explanatory. Involuntary eye movements are in many cases reflexive; driven by head movements or stimulation in the environment. The large scale eye movements are ones that are observable by the naked eye. The miniature eye movements are too small to be observed without some form of magnification. Some of these eye movements are miniature eye movements are interesting, as will be seen in the next chapter, because they always keep the eyes moving. Our eyes are never completely still, and this will have some profound implications. The version eye movement. In the vergence eye movements, the eyes move in parallel. The make the same size and direction of movement. In the vergence eye movements, the eyes move in the opposite direction. In this chapter, we will discussion voluntary large scale eye movements of both version and vergence types.

#### Version Eye Movements

There are two main types of voluntary version eye movements: **Saccades [to glossary]** and **Smooth Pursuit [to glossary]**. Each type of eye movement will be discussed in turn.

**Eye Muscles.** To understand how we move our eyes, let us take a look at some of the most active muscles in our body, the eye muscles. Open **Interactive Illustration 4.x**, **Eye Muscles [link to media NOT DONE]** which shows the eye muscles are placed on the eye. We have six eye muscles on each eye that are used to move our eyes. Four of the muscles attach at close to right angles to the eye and are called the rectus muscles. They are: the superior rectus, attaching on top of the eye; the inferior rectus, attaching on the bottom of the eye; the medial rectus, attaching on the side of the eye near to the nose; and finally the lateral rectus, attaching on the side of the eye nearest to the outside of the head. These muscles move the eye up and down (superior and inferior) and from side to side (medial and lateral). You can see these eye movements simulated in the model. Just press the button for the appropriate eye movement, and the illustration will show which muscle contracts and which relaxes to make each of these eye movements.

The other two eye muscles are the superior and inferior oblique muscles. They come in from odd angles and serve to rotate the eyes around the horizontal axis to a small degree. If you tilt your head to the side, your eyes may rotate a small degree in the opposite direction to stay vertical. They have a limited range of movement in this direction, only about 10-20 degrees (REF).

**Saccadic Eye Movements.** Saccadic eye movements are our most common type of eye movement (REF). They are very rapid, with the eyes spinning at speeds of up to 400 deg/sec. We use these eye movements to look around from one object to another. Their main purpose is to bring an object to the foveas of both eyes, so that we can use our foveas to investigate the object. There are several

features of these eye movements that make them interesting. First, they do not happen immediately when you desire to make the eye movement. While the delay is variable and depends on many factors, a delay of about  $1/5^{\text{th}}$  of a second (200 msec) seems to work as an average (REF).

Open Interactive Illustration 4.x, Eye Movements to see an illustration of these eye movements. When you open the illustration, two eyes will be drawn on the right side of the screen. In the middle of the screen are 5 dots which are 5 different objects that the eyes in this illustration can look at. Green lines connect the dot being examined to each of the eyes' fovea. The eyes are looking at the middle figure which is about in the middle of the screen. Clicking on each of the objects will cause the eyes to look at that object. The change in which object is being examined happens very quickly but you might notice a delay. Not all of that delay is due to slowness of your computer. Some of that delay is actually programmed in to simulate the delay from the beginning of the intension to make a saccade to the actual beginning of the saccade.

Open Interactive Illustration 4.x: Saccadic Eye Movements [link to media] to experience this delay in your own eye movements. This is a very simple demonstration. When you press the **Start** button, you will be tracking a small circle back and forth across the screen. When the dot is on the left side of the screen, it is at its home position and will reset there for approximately a second. Then it will move to some location on the right half of the screen (randomly selected). You are to move your eyes to this new location. And then, when the dot returns to the home position on the left, move your eyes back. Some times the dot will return to its original spot after only a very short period of time. See if you ever notice that it returns to its starting time sometime before you make your eye movement.

Also see if you ever found yourself trying to not make an eye movement, because of the return to the start, but your eyes actually made the eye movements before you could stop them. Under normal circumstances, these saccades reach a point in the neural computations to make the eye movements so that they tend to go forward regardless if the reason for the eye movement still exists or not. This characteristic of saccadic eye movements is called ballistic, like ballistic missiles. Once you launch them, there is no calling them back.

Saccades and Heads Up Displays (HUDS). There are some safety issues related to saccades that human factors and other engineers have struggled to overcome. When you are looking down at your speedometer or your radio in your car, you are not looking at the road. Now during the delay in making saccades, first down and than back up; it is possible you could miss some important event. In a car, the delay is usually not a big problem which is why most of us are still alive; but in a military jet, that delay of looking down to the instrument panel could be fatal. To get around this problem, in military aircraft they project a version of the instrument panel right on the cockpit in a transparent form that can be seen through. This allows the pilots to read displays without looking down. You might have seen HUD speedometers in some cars. In Chapter 8 on Depth Perception, we will see some reasons why this technology has not been safely transferred to cars.

Saccadic Suppression. Saccadic eve movements are very fast. They last only a few milliseconds and can cover a large portion of the world in front of us. They move much faster than our temporal integration period. This is like having a camera open too long, and having the images get blurry like the picture in Figure 4.x. The young man in the background is moving while the camera shutter is open and looks like a blur. The same would happen, if he sat still and the camera swept across the seen, only the entire picture would be blurred. However, you rarely have any experience of this blur because of a set of mechanisms that collectively are called saccadic suppression [to glossary] (Latour, 1962; Volkmann, 1962; Volkmann, Schick, & Riggs, 1968). Saccadic suppression causes us to not notice, under normal circumstances, the effects of the blur caused by saccades, or even that we make nearly as many saccades as we do. I want to challenge you to do some observational experiments. For the first experiment, go to the mirror and while you hold your head very still, look from one eve to the other and see if you can see your eyes move. If you hold your head very still, you will not see the movement. For the second experiment you will need a partner and something to read. First, read and see how you experience moving your eyes across the page. Second, while your partner holds the reading material in a position you can observe, see if you can observe how they move their eyes. Most of us describe our eye movements as smooth most of the time while reading. What you will see is that the eye movements are the jerky saccades we have been describing.

The suppression of the blurry images seems to be due to many factors including the blur itself, the masking or suppression of the blur by the clear images after the eye movement, and active suppression of

the visual system during saccadic eye movements (Volkmann, 1986). Let me discuss the masking mechanism a little bit, as it an important visual phenomenon in its own right and it is the least intuitive. In perceptual masking, one stimulus makes it harder to perceive another stimulus.

In masking associated with saccadic eye movements, we are dealing with a situation where the masking stimulus, the clear scene at the end of the eye movement, masks a stimulus that occurred earlier, the blur of the saccade. This type of masking is called **backward masking [to glossary]** (Breitmeyer, 1984). Figure 4.x gives an overview of a backward masking experiment. The target stimulus in a standard experiment is a small circle. It is this stimulus that the participant is trying to see. It is presented first for a brief period of time, and then the masking stimulus is presented. Notice its rather odd shape, it is an annulus stimulus that surrounds and just touches the edge of the target stimulus. The masker is not physically on top of the stimulus it is trying to hide. This situation works great for the use of backward masking in saccadic suppression. In fact, the blurs will move right up to the edge of the clear images, putting the blurs in a great location to be masked. In fact, the way backward masking works makes it so useful for saccadic suppression that it has been argued that the reason backward masking exists in the visual system is to support saccadic suppression (Matin, 1974).

Let us try a backward masking experiment. Open **Experiment 4.x, Backward Masking [link to media]**. The first page is the stimulus parameter page. You can adjust many parameters of the experiment such as the duration of the fixation before the stimuli are presented, the time between the removal of the fixation stimulus and the presentation of the target stimulus, the duration and size of the target stimulus, the time between the target stimulus and the masking stimulus, the space or gap between the target stimulus and the size and duration of the masking stimulus. For now leave all values at their default levels and pres the **Done** button. Next is the psychophysical method menu which is for a Method of Constant Stimuli method. Pressing the **Done** button will start the experiment.

In each trial, a fixation mark will be presented. Stare at it. Then the target stimulus will be presented followed by the masking stimulus which surrounds the target stimulus. After each trial you will be asked "**Did you see the stimulus?"** at the bottom of the screen. Press the **Yes** button if you did see the target stimulus, and press the **No** button if you did not see the stimulus. After this condition is done, the experiment will be repeated without the masking stimulus. At the end of the experiment, the thresholds for the target stimulus with and without the masking stimulus will be presented. Do the experiment and then return to the text.

If your results match what is normally found, the threshold for the **With Masker** condition should be significantly higher than for the **No Masker** condition. This finding indicates that it is harder to see the target stimulus when the masker both follows and surround the target. This is a good example of backward masking.

[Can I do a suppression experiment?]

**Smooth Pursuit Eye Movements.** The next type of version eye movement that we will discuss is the smooth pursuit eye movement. These eye movements are well described by their name. They appear very smooth in comparison to the saccades, and they are used primarily in tracking. You need to have a stimulus to follow with these eye movements or you will be making saccades. Reopen <u>Interactive</u> <u>Illustration 4.x, Eye Movements</u> to see a simulation of smooth pursuit eye movements. This time you will be using the slider on the left hand side of the screen. When you move that slider up and down, the fixation objects will move up and down and the eyes will try to track the object that they have been examining. Move the slider quite slowly at first up and down. As you do this you will observe that the eyes will rotate and keep the fixation dot centered on their fovea. However, smooth pursuit eye movements do not go very fast; they have a top speed of about 10 to 20 deg/second (Alpern, 1982). When the motion of the object exceeds this period of time, then you will see saccades begin to appear in the eye movements as they eyes use the saccades to try to keep up with the moving object. If you move the slider faster you can observe that the eyes will begin to fall behind and saccades will occur that catch the eye up, for a moment.

Now let us try a little, informal, experiment to experience a few of these factors about these eye movements. You will probably need a partner to observe your eye movements, as it is not always clear when we change from smooth pursuit eye movements to saccades. Recall the saccadic suppression discussed above. Open **Experiment 4.x: Smooth Pursuit Eye Movements [link to media]**. This will be a very simple experiment. First the screen is black. Holding your head up so your partner can observe

your eyes, see if you can move your eyes across the screen smoothly without using saccades, even very small ones. You will not be able to. Then press the **Start** button at the bottom of the screen. A dot will appear and start moving across the screen. Follow this dot. It will move back and forth across the screen in a fairly smooth motion (it may stop occasionally, that is a result of using Java to write these routines). The slider to the right of the **Start** and **Stop** button controls the speed of the dot. Start with a slow speed where you can easily keep up with the dot. Have your partner verify that your eyes are moving smoothly. These are smooth pursuit eye movements. Then gradually increase the speed of the dot, and with your partner, you will find a speed of the dot where you will be unable to keep up with the dot using smooth pursuit but will start using saccades to keep up. This speed will occur at the upper speed of smooth pursuit eye movements. Your partner will see your eyes start jumping across the screen. You can also trade places with your partner so that you can observe these changes in eye movements. Go do the task at this point in time and then return to the book.

Now that you have performed this task, you and your partner might have noticed several observations. First, smooth pursuit eye movements take a while to start, so there is usually a beginning saccadic eye movement to start the tracking. This can be observed by repeatedly starting and stopping the dot motion. Recall that smooth pursuit eye movements are not very fast. So as you increased the speed of the dot motion, saccades crept into your tracking of the object. Okay, now think about the batter in baseball. The batter is coached, and I have told kids I have coached to do this, that the batter must follow the baseball with their eyes from the pitcher all they way in to where the ball makes contact with the the bat. The ball is moving a lot faster than smooth pursuit can keep up. While some of the ability to track a pitch is helped by reflexive eye movements that help tracking faster stimuli, it comes down to the fact that batters cannot track a ball all the way in (Bahill & LaRitz, 1984). That does not mean a batter should not try to tack the ball; however, it is apparent that they cannot be completely successful.

#### Vergence Eye Movements

The vergence eye movements are when the eyes either rotate towards or away from each others. These eye movements are used to be able to examine objects at different distances. To see an object clearly, it is not sufficient to have the image of that object fall only on the fovea of one eye. It is necessary to have the image of that object fall on the foveas of both eyes. However, the eyes are in different locations and, thus, have a slightly different view of the world. To get an object on both foveas, they need to rotate relative to each other, and how much rotation is needed depends upon the distance of the object. Reopen **Interactive Illustration 4.x, Eye Movements [link to media]** to see these eye movements in action.

In this case, you will be using the slider at the bottom of the screen. When you move this slider, the objects will move closer or farther from the eyes. What you need to observe in this case is how the eyes now rotate relative to each other to try to keep the images of the object they are looking at on the fovea.

First, let us move the object to the extreme distance on the left allowed on the screen. As you do, the eyes move apart. This rotation of the eyes apart from each other is called **divergence [to glossary]**. Now move the object as close as possible to the eyes on the right side of the screen. Notice how the eyes rotate together. This rotation together is called **convergence [to glossary]**. You can also move the object from side to side. As you move the object from left to right and back again, notice that the eyes rotate together. This would be a version eye movement and since it tracks, this simulation would be a simulation of smooth pursuit eye movements.

The vergence eye movements actually are controlled much like saccadic eye movements. They take a fairly long time to start and work in a moderately ballistic fashion. This has led some researchers to hypothesize that there might be some suppression during vergence eye movements much the same way that there is during saccades. This conclusion was subsequently supported (Manning & Riggs, 1984).

#### **Summary**

This chapter has covered a lot of ground. The chapter started with a discussion of the absolute threshold. To determine the absolute threshold for vision, a lot of fundamental visual factors had to be considered. These were the state of adaptation, the size and duration of the stimulus, the wavelength of the stimulus, and the location of the stimulus in the eye. These factors relate to dark/light adaptation, spatial and temporal summation, spectral sensitivity, and the arrangements of rods and cones in the retina. Dark and light adaptation refers to the fact that we need to adjust our sensitivity to the light level of the surroundings. Dark adaptation proceeds more slowly than light adaptation. One of the major factors in

dark and light adaptation is the change from the cones, used in the daytime, to the rods used at night and vice versa. The differences are so great between the ways that rods and cones work that researchers have described the eye as having two types of vision, photopic and scotopic. Photopic refers to daytime vision and scotopic to nighttime vision. Spatial summation refers to the fact that to increase our ability to detect light, we add up the light that falls on a small region of the retina. Temporal summation is similar but the light is added over time. Spectral sensitivity refers to the discovery that rods and cones are not equally sensitive to all wavelengths. Moreover, rods are relatively more sensitive to shore wavelengths and cones to long wavelengths. A difference in relative sensitivity is known as the Purkinje Effect. Cones are concentrated near the fovea. There are no rods in the fovea, and their greatest concentration is about 20 degrees into the periphery.

After all of these factors in vision were discussed, the chapter covered acuity. There are several types of acuity. In standard acuity, we can resolve features down to about 1 arc minute. However, for a positional acuity, vernier acuity, we can resolve differences down to about 5-7 arc seconds. Acuity is greatest in the fovea and falls off rapidly in the periphery.

Finally, the chapter discussed eye movements. Eye movements are accomplished by the six eye muscles of each eye. There are two types of voluntary eye movements: version and vergence. In the version eye movements, we move our eyes in the same direction. The most common version eye movements are saccades, where we look around from object to object. Saccades move so fast that the visual system actually makes us insensitive to stimulation during saccades. We track objects with smooth pursuit eye movements. In the vergence eye movements, the eyes either rotate toward or away from each other. They are used primarily to look at objects at different distances.

With this information as background, we are ready to start discussing how the visual system works to take in stimuli and make a pattern out of the stimulation. This effort is central to our perception of objects as objects, and it is called form perception.

#### Problems

- 1. Explain why tail lights are red how does that color help the driver following see at night
- 2. Why are the blue headlights that are one some cars a really bad idea
- 3. What is different about looking at red LED clock lights at night and the blue LCD lights at night why?
- 4. Take a red object and a blue object outside say to point. View as dusk falls to twilight what changes happen in them and why?
- 5. Why can you write out words with sparklers at night?
- 6. Why can't you watch a regular television outside during the day?
- 7. What are crucial factors in deciding the size of dots that make up computer screens?
- 8. Is it possible to see wheels rotating backwards in the real world?
- 9. What conditions make it hard to see electronic displays in cars and why?

Structure	Tons	Pounds	Scientific Notation
The Battleship Missouri	58,000	116,000,000	1.16x10 <sup>8</sup>
The Empire State	365,000	730,000,000	7.3*10 <sup>8</sup>
Building			

Table 4.x. Conditions for Measuring the Absolute Threshold

Condition	Value	
State of the Eye	Dark Adapted	
Location of the Stimulus	20 deg in the temporal retina	
Duration	Longer than temporal summation	
Size	Larger than spatial summation	
Wavelength	At peak sensitivity for rods	

Table 4.x . Photopic vs. Scotopic Vision

	Photopic	Scotopic
Color Vision	Yes (3 cone classes)	No (Only one type of rod)
Peak Spectral Sensitivity	~ 550	~505 nm
Acuity	Good	Poor
Temporal resolution	Good	Relatively poor
Regions of Greatest Sensitivity	Fovea	~20 deg in Periphery
Relative Sensitivity	1,000,000	1



Figure 4.x. Effects of moving a camera while the shutter is open. On the left, (a) the shutter is open briefly and the image is clear. On the right, (b) the image is open for a longer period and the camera is moved. The image is blurry.





Figure 4.x. A graph comparing the relative spectral sensitivity of rods and cones.





Figure 4.x The outline of a backward masking experiment.