Chapter 10

Fundamentals of Audition

Chapter Outline:

I. Introduction
II. What is the Auditory Stimulus?
  a. Basic Dimensions of the Sound Stimulus
  b. Phase
  c. Timbre
III. What are all the Parts of the Ear and How do They Work?
  a. The Outer Ear
  b. The Middle Ear
  c. The Inner Ear
  d. What does the Sound Stimulus do to the Inner Ear?
  e. Review of the Auditory Stimulus in the Ear
  f. Auditory Transduction
    i. Ohm's Acoustical Law
    ii. Fourier Analysis in Audition
  g. Auditory Pathways
IV. Fundamental Dimensions of Sound
  a. Encoding Intensity
    i. Auditory Thresholds
    ii. Equal Loudness Contours
  b. Encoding Frequency
    i. Frequency Discrimination
    ii. Place Theory
    iii. Frequency Theory
    iv. Reconciliation
  c. The Effect of Timbre
V. Summary

Fundamentals of Audition

It is a classic and overwrought piece of melodrama often used in movies. The hero is sneaking up on the movie villain and steps on a stick or bumps into something and instantly gives away the secret and the location of the attack. It wouldn't be fair for the good guy to use intelligence to advantage and would certainly shorten the climax of the movie. This circumstance reveals one great difference between audition and vision as senses. Vision is restricted in us to what we can see in front of us. We have a limited field of view. If we are discussion what we can see in detail, the restriction is much smaller and tied to the fovea. We hear in all directions. Thus, this is one advantage that hearing has for us.

Still, this does not really tell us what is hearing. The simple answer is that hearing is the response to sound stimuli in our environment. More deeply this question begins to ask what is a sense? While this question will be dealt with even more deeply in the somatorsensory system chapter, there are some unique aspects of this question that should be addressed here. Certainly to hear responds to a very different type of stimulus than seeing does. The world, as represented by sound, is very different than that experienced by sight. Yes, sound does and can represent the world, think of bats and owns who use sound, both that they produce and naturally in the environment, to fly and catch prey (REF). Definitely, these animals have a rich and detailed experience of the world that is based not on sight but on sound. So, hearing, in one sense, is our experience of the world via sound stimuli which is different than just saying our experience of a sound stimulus. Now the stimulus has a significance and importance. To illustrate, open **Experiment 10.x, Sound Environments [link to media]**. In this interactive illustration you get to play three random sounds [FIGURE OUT THE INTERFACE AND DESCRIBE IT HERE]. In each of the X cases you can play the same sound, a set of foot steps walking, and then enter what you can figure out about the environment they person is in as they are walking. After you play each sound you will be given a chance to select from a drop down menu of a list of options [NOTE TO ME HAVE MORE OPTIONS ON THE LIST THAN ARE USED]. See if you can identify the environment that the recording takes place.

If you are like most listeners you should be very successful at this task. The ability to successfully identify the environment indicates that sound does convey a great deal of information about our environment. This information is not the same information about the environment as is provided by vision. Audition enhances and complements the experience of the world that is experience by vision and often we do not take the time to separate the two sources of information about our world. So a world without sound is a world that is limited and our knowledge of the world is limited without sound.

In this chapter and the next chapter, this text will explore the nature of audition. Recall the model of sensory systems that were provided in Chapter 1. This path will be followed to help develop an understanding of the nature of audition. Thus, the stimulus will be discussed first, followed by transduction and the pathway that auditory stimuli follow from the ear to the cortex of the brain. Then, simpler aspects of our auditor experience will be discussed followed in the next chapter by more complex auditory experiences.

What is the Auditory Stimulus?

At our ear, the auditory stimulus is periodic variations in air pressure. Over time, air pressure will increase and decrease slightly and these small changes in air pressure constitute sound to our ear if they occur strongly enough and fast enough but not to fast. To make the nature of the sound stimulus clearer open **Interactive Illustration 10.x, Sound Basics [link to media]**. This illustration will allow you to interact with a sound stimulus and understand some of its characteristics. To start the illustration, click on the **Reset** button at the bottom of the screen below the graph. This graph will now show a sine wave much like the illustration of the stimulus for vision. The y-axis of the graph is air pressure and the x-axis of the graph is time. Since sound is varying in time in pressure, the stimulus for sound is often described as a sine wave just like the visual stimulus. You can hear the sound created by this sine wave if you have a sound card and speakers on your computer by clicking the **Play** button at the bottom left corner of the screen. You can stop the sound by clicking the **Stop** button next to the **Play** button.

Since the sound stimulus is described as a sine wave, it is made up of a repeating stimulus with each repetition called a cycle. You can show one cycle of the sound stimulus by clicking the **Show First Cycle** check box on the lower left corner of the screen. The first cycle will then be drawn in blue to discriminate it from the rest of the sound stimulus. The first half of the cycle is where the air pressure is greater than the average air pressure of the air that the sound stimulus is traveling through. This increase in air pressure is where the air density has increased slightly is called a **compression [to glossary]**. You can see the first compression highlighted by clicking the **Show First Compression** check box. The first compression will be drawn in red. Selecting this option will clear the highlighting of the first cycle if it is still on. The second half of the cycle where the air pressure is less than the average air pressure of the air that he sound stimulus is traveling through. This decrease in the air pressure is where the air density had decreased slightly is called a **rarefaction [to glossary]**. You can see the first rarefaction highlighted by clicking the **Show First Rarefaction** checkbox. The first rarefaction will be drawn in magenta. The options to show the first cycle, compression and rarefaction will clear the other options.

Basic Dimensions of the Sound Stimulus

If you recall, when measuring the stimulus for light, there are two important measures of a sine wave, the amplitude and the wavelength. These measures apply to the sound stimulus as well. The amplitude or intensity of the sound is most closely associated with the **loudness [to glossary]**. You can experiment with how intensity relates to loudness by pressing the **Play** button and then adjusting the **Intensity** slider on the left side of the screen. When you start the program, the intensity is at its maximum as you slide the **Intensity** slider down, the height of the sound wave and the amplitude of the sound both decrease. Readjust the intensity to a comfortable level.

Intensity in sound is most commonly measured in decibels (abbreviated dB) in sensation and perception. Decibels are 1/10th of a Bel which is a unit of sound intensity named in honor of Alexander Graham Bell. The equation for the dB is:

(1)

or

(2)

$$dB = 20\log\left(\frac{P_1}{P_r}\right)$$

Both equations will gi There are two important features of this equation to not which is why you are being tort t, actually sound pressure amplitude is P_1 in the equations. The decibel equation actua nsity difference between this intensity and another intensity, P_r. So, sound measures of decibels are measure of relative intensity and so it is important to know what both sounds are. Sometimes there is only one sound that is of interest. In this case, there is an implied P_r which is near auditory threshold that is being used as the comparison sound. In this case, the measure should be give as dB_{spl}, where the SPL stands for sound pressure level.

The other major measure of a sound stimulus is related to the wavelength. In audition, it is more typical to talk about the **frequency [to glossary]** of the sound stimulus, not its wavelength. The frequency a sign wave, including a sound stimulus, is the number of cycles that occur in a period of time, usually a second. Remember, that each repetition of a sound stimulus is one cycle. An easy way to think about frequency is how many peaks occur in a second. When cycles per second is used as the measure of frequency, the frequency unit is called **Hertz** (**Hz**) **[to glossary]**. So a sound stimulus that has twenty peaks hit your ear in a second would have a frequency of 20 Hz. Frequency is inversely related to wavelength. As frequency gets higher in value, the wavelength gets longer. The frequency of a sound stimulus is most closely related to the **pitch [to glossary]** of a sound. Pitch refers to the experience of whether a sound has a high or low quality to it. For example a soprano has a high voice and a base has a low voice. The keys on the right hand side of a piano play high pitches and the keys on the left hand side of a piano have low pitches. You can see how frequency relates to pitch by adjusting the **Frequency** slider on the extreme left side of the screen. As you lower the **Frequency** slider, the wave lengthens which will also lower its frequency and the sound has a lower pitch. The situation is reversed for moving the **Frequency** slider up. The normal health young adult can hear sounds across a range from approximately 20 Hz to 20,000 Hz (also shown as 20 KHz for kilohertz).

Phase

In addition to these basic characteristics of a wave form that need to be considered, a couple of additional features become very important in audition. The first is the **phase [to glossary]** and the second is the collection of frequencies present in the sound. Phase refers to the position in a sound wave. Looking at **Interactive Illustration 10.x, Sound Basics** will help illustrate this point. The sound stimulus is a repeating pattern. One repetition has 360 degrees in it starting at the far left side of the waveform. At this position air pressure of the sound stimulus is at the mean level and it is at a phase of 0 degrees. The peak is

1/4th the way through the cycle of the wave, and is thus at 90 degrees of phase. When the stimulus returns to the mean level, it is at 180 degrees of phase. The tough, or rarefaction, is 3/4th the way through the cycle and is at a phase of 270 degrees of phase. The cycle is complete when it reaches the mean pressure again, and it is at 360 degrees of phase for the first cycle, but 0 degrees of phase for the next cycle. You can see one cycle of the sound stimulus highlighted by clicking on the **Show First Cycle** check box in the lower left corner of the screen. Now the first cycle of the sound stimulus is shown in red. You can manipulate the frequency and intensity of the stimulus to see how the first cycle changes with these two variables. When you click the **Show Phase** checkbox right above the **Show First Cycle** checkbox, moving the mouse over the graph of the sound way will help you see the phase for each position of the sound stimulus. A line going from the top of the graph to the bottom of the graph will appear as well. This gray line will intersect the sound wave at the point being indicated. Near the tip of the mouse, there will be text indicated the phase for the cycle at that position on the graph. Start at the left hand edge of the graph, and the phase will be 0. Track along the first cycle, and you will see the phase increase. At the first compression, the phase will be 90 as indicated above. When you reach the end of the first cycle, the 0hase returns to 0 and the values for phase start over, just as the cycle does. The reason that phase is important to understand is that it affects what we hear.

Open **Interactive Illustration 10.x, Phase and Cancellation** **[link to media]**. In this figure, there are two graphs of sound waves. The top is green and the bottom appears red. While the two graphs are very similar, they are actually quite dissimilar. In this example, the computer is actually generating two sound waves and the bottom graph is drawing both of the. One sound wave is drawn in red and the other sound wave is drawn in cyan. They are exactly the same and since the red waveform is drawn second, it does not show yet. The top graph is the combination of the two lower graphs. It is the sound stimulus that actually hits your ear. It looks very similar to the bottom graph because since the two sound waves shown in the bottom graph are identical, the combination of the two graphs is a sound stimulus of the same frequency, but with twice the amplitude. If you click the **Play** button you will hear a single tone. In fact the ear would have no way of knowing that two tones are present.

If you use the **Phase** slider on the right hand side of the screen you can adjust the point in the cycle of the red waveform. As you change the relative phase of the two waveforms, the first thing you will notice is that you can now see the cyan wave form, but there is no appreciable change in the summation of the two wave forms. Now, adjust the **Phase** slider to 180 deg. You can see what the starting phase of the red waveform is by reading the value at the bottom of the phase slider. Now what you see are the red and cyan waveforms completely opposite each other. Where the red waveform goes up, the cyan goes down and vice versa. The top graph is the point by point summation of those two waveforms and you can see that they cancel each other out. So the resulting waveform, even though made up of two waveforms, is flat. This is known as **cancellation [to glossary]**. When two sounds of the same frequency have a phase difference of 180 deg they cancel each other out. When the phase difference is 0 deg (or 360 deg) the two waveforms complement each other. The significance of this feature of sound waves is that the world is composed of many sounds each of with usually is composed of many frequencies. These sound waves interact at the ear and can either add to each other or even cancel each other out.

However, most stimuli are not of the same frequency. In this case, the situation gets a little bit more complex. Use the same interactive figure but return the **Phase** slider to 0 deg. You will now be manipulating the **Freq (Hz)** slider which will change the frequency of the waveform drawn in red in the lower graph. Click on the **Freq (Hz)** slider but do not move it. This action selects the slider and now your arrow keys will move the slider one step at a time and for the first few manipulations of the slider, it will be nice to have that fine of control.

Move the slider so that the frequency of the red curve, indicated by the **Freq:** at the top of the top graph is 501 Hz. The cyan sound wave will always be 500 Hz. Now the red wave has a just a bit higher frequency than the cyan wave. Since the phase is 0, they start out in synchrony but gradually the red wave will get ahead of the cyan wave and they will become out of sync. You will see this on the lower graph by seeing the cyan waveform peaking out behind the red wave form to the right side of the graph. You probably cannot tell but the height of the combined waveform gets a bit smaller as it goes to the right because the waveforms are not completely in sync. You can see the overall patterns by clicking on the **Full** button at the bottom right of the top graph. This button will plot one full second's worth of the data on the screen. You will not be able to see the individual peaks and valleys but you will see the overall shape of the waveform. As you will be able to see, the waveform gets smaller until there is no amplitude in the middle of the second and then it returns to full size at the end of the second. Press the **Play** button to hear what this stimulus sounds like. Instead of two frequencies, it sounds like one tone that gets softer and then louder. The rate of this loudness changes is one time every second as you can see from the graph. This changing loudness is called a **beat [to glossary]**. When two tones are close together in frequency are played at the same time, you hear beating and the rate of the beating is equal to the difference between the two stimuli in Hertz. Move the **Freq (Hz)** slider to 502 Hz and see that there are now two places where the amplitude of the combine wave goes to 0. If you are still listening to the resulting sound you will hear the rate of beating double. At this point it might be a good idea to hit the **Reset** button on the top graph. As you increase the frequency, slowly, not using the slider, you will hear the rate of beating keep increasing until it gets double. You will even begin to see the beating appear on the now zoomed in top graph and you will see the two waveforms on the lower graph go into and out synch. Eventually as the two frequencies get farther apart, the beating becomes a roughness and then you will hear two separate tones. Eventually the two tones are far enough apart in frequency, the beats are not perceptible and instead, you hear two tones.

Timbre

The last dimension of the sound stimulus to be considered is **timbre [to glossary]**. Most sound stimuli are composed of multiple frequencies of sound and most of those frequencies do not cause beating. One of the ways the presence of multiple frequencies impacts our hearing is in the perception of timbre or the quality of the sound. In a musical sound, while many frequencies are present, there is one frequency that determines the pitch that is experienced and that is called the fundamental frequency. The other frequencies are called harmonics or overtones. A flute and a clarinet can play the same note but you will rarely confuse the two instruments. [RECORD FLUTE AND CLARINET AND DO AN FFT ON THEM AND DEVELOP AND ACTIVITY] Both instruments produce multiple frequencies and when playing the same pitch the fundamental frequency is the same. However, the harmonics differ in terms of strengths and sometimes which ones are present. This pattern of differences leads to the different qualities of sounds even when producing the same pitch.

Open **Interactive Illustration 10.x, Timbre** **[link to media]** to explore the nature of timbre. When the interactivity opens, there is a single frequency present in the stimulus and you can see a single sine wave on the graph of the sound stimulus at the top of the screen. You can press the **Play** button in the lower left corner of the screen to start the sound. You will hear the sound indicated on the graph. In the middle of the screen are 10 sliders numbered from **0** to **9** with two buttons next to the sliders, the **Zero All** and **Zero Harmonics** buttons. With these sliders and buttons you can control what frequencies are present in the stimulus. The fundamental is the **0** frequency on the screen. The frequency of this ton is 200 Hz which is a fairly low frequency. The frequency of A above middle C is 440 Hz for comparison, so this tone is more than 1 octave below that A. Each of the other harmonics is a whole multiple of this frequency. Thus, the **1** harmonic is 400 Hz, the **2** harmonic is 600, the **3** harmonic is 800, etc. The height of the slider indicates the amplitude or intensity of that harmonic. At the start the **0** harmonic is about at half of its loudest possible intensity and the slider is about in the middle of its ranger. All of the other harmonics have an intensity of zero and their sliders are at the bottom of their ranges.

Click on the **Play** button while only the fundamental is present in the sound stimulus. While listening sound, start adding new frequencies. Add them at a very low intensity so that you don't create a sound too loud for your computer to generate. If you start seeing flat areas at either the top or the bottom of the graph of the sound stimulus, reduce the intensity of your fundamental and/or harmonics so that it again looks curved. As you add the new frequencies, particularly at this low intensity, you will probably not hear the separate pitch (unless you are an experienced musician or with very good hearing) but one tone and the quality of that tone will change. To me, when I add several of the harmonics, starting from the lowest harmonic, **1**, it sounds a bit like a cheap electric organ. There is more to timbre than has been discussed but this discussion gives a feel for the topic.

Resonance. [Add a section on resonance.]

[NEED TO RECORD SOME SOUNDS TO ADD TO THIS INTERACTIVITY]

What are all the Parts of the Ear and How do They Work

The sound stimulus enters the head at the sides of our sculls where we have two fleshy protrusions that in common parlance we call our ears. As in vision, without transduction of the auditory stimulus, there would be no hearing. So the question for this section is how does transduction take place? We will start with the entrance of the sound stimulus into the side of our heads and follow it until transduction occurs.

To give you an overview of the path that will be followed, see Figure 10.x. As the path of the sound stimulus is explored, it will be discovered that the ear can be divided into three main sections. These sections are the outer ear, the middle ear and the inner ear.

The Outer Ear

The outer ear begins with the **pinna [to glossary]** which is what we usually refer to as the ear. It is the fleshy extrusion on the side of your head. The pinna joins the **external auditory canal [to glossary]**. If you look at Figure 10.x, you might see that the pinna plus the auditory canal look a bit like a funnel and that is a good description of their function. They gather sound waves and channel them to the ear. The analogy should not be a great surprise and sound travels through the air which is a gas and gases are fluids just the liquids you put through the funnels in a kitchen or to guide oil into a car engine. The **tympanic membrane [to glossary]**, more commonly known as the ear drum, seals the end of the outer ear. The ear drum is a good term as the molecules of the air in the sound stimulus will beat upon the tympanic membrane to move the sound stimulus to the next part of the ear.

The Middle Ear

So, to this point, the sound stimulus is pressing against the tympanic membrane and moving it back and forth. The next question is what is on the other side of the tympanic membrane and what is going on in this area. The space on the other side of the tympanic membrane is the middle ear. The middle ear is an air filled space in your skull. A drawing of this space is shown in Figure 10.x. The primary active elements of the middle ear are the three smallest bones in your body, the **ossicles [to glossary]**. The names of the ossicles are, in order from the outer ear to the inner ear, the **malleus [to glossary]** (hammer), **incus [to glossary]** (anvil), and **stapes [to glossary]** (stirrup). These bones transmit the sound stimulus from the tympanic membrane to the inner ear. **Interactive Illustration 10.x, The Middle Ear** shows the three ossicles with the malleus attached to the tympanic membrane and the stapes connected to the oval window which is part of the inner ear. Pressing the **Start Button** will start the tympanic membrane moving and the illustration will show how the ossicles move in sympathy with the tympanic membrane. At the end of the ossicles, the stapes moves in a piston like fashion reproducing the sound stimulus against the inner ear.

There is one other structure in the middle ear that is important to our hearing. That is the **Eustachian tube [to glossary]**. From Figure 10.x you can see the Eustachian tube entering the lower part of the middle ear. The other end of the Eustachian tube connects to the back of the throat. Normally, the Eustachian tube is closed but it opens briefly when we swallow or yawn, for example. The brief opening of the Eustachian tube is connected to what is known as popping of your ears. The middle ear is filled with air, and except for when the Eustachian tube is open, it is cut off from any changes in air pressure of the environment we are in. For optimal operation of the tympanic membrane, it is important for the air pressure on both sides of the tympanic membrane to be equal. The air pressure outside changes due to weather and altitude. The Eustachian tube opens to allow the two air pressures to equalize. Usually the change in air pressure in the middle ear is slight which is why it does not open very often. Yet, when you ascend or descend a mountain, there can be a large change in air pressure. If you are ascending the mountain, the air pressure on the outside of your ear decreases with each step up in altitude. If your Eustachian tube stays closed, you will have much more air pressure in the middle ear side of your tympanic membrane and outside of it causing it to bow towards the outside of your head. Being pushed out makes so that it does not move with sound stimuli as well and in these cases you might find your hearing is not quite as good as normal. If you yawn or swallow in these cases there will be a rapid reduction of the air pressure inside of your middle ear which you will experience as the popping of your ears.

At this point is might be useful to review travel of the sound stimulus. First, the sound stimulus encounters the pinna which helps to funnel/direct the sound stimulus into the external auditory canal where it comes in contact with the tympanic membrane causing it to move in sympathy with the sound stimulus. When the stimulus moves towards compression it pushes the tympanic membrane in and when the sound stimulus moves towards rarefaction, the tympanic membrane is pushed out by the pressure of the air in the middle ear. Thus, the sound stimulus is transferred into the motion of the tympanic membrane. The motion of the tympanic membrane is transferred to the malleus which moves the incus and causes the stapes to move in a piston like motion in synchrony with the motion of the tympanic membrane. At this point the sound stimulus is now transferred to the inner ear.

The Inner Ear

The stapes does not do transduction. Examination of the stapes shows that it is merely pounding against a membrane. From this observation, it can be concluded that there is a further step between the reception of the sound and transduction. This membrane can be used to divide the ear into another section called the inner ear. Scraping the inner ear away from the rest of the head, it can be observed that the inner ear is contained in a single bony structure called the **cochlea [to glossary]**. Figure 10.x shows how the cochlea appears in the head. To learn more about the cochlea it would be nice to open it. **Interactive Illustration 10.x, The Cochlea** unrolls the cochlea and opens it to allow you to see some of the internal structures of the cochlea more clearly. The cochlea gets its name from its snail-like shape. Cochlea is Latin for snail or being in the form of a snail shell.

Opening the cochlea reveals a great different between the inner ear and the two other sections of the ear. The inner ear is fluid filled. The outside is covered with a hard bony surface, so for the stapes to have much of an impact on the cochlea there needs to be a soft area. This soft area is that membrane mentioned above under the foot of the stapes. This membrane is called the **oval window [to glossary]**.

However, unlike air, water does not compress very much. If the stapes is to be able to push in on the oval window, there has to be some place for the fluid to move. There must be another soft place where the fluid can push out. This place for the fluid to move is provided by the **round window [to glossary]** which is below the round window on the cochlea, but at the other end of the flow of the fluid in the cochlea.

Let us look more at **Interactive Illustration 10.x, The Cochlea**. The place that the stapes contacts the cochlea is shown on the left side of the image. The stapes is up against the oval window and the round window is indicated below the oval window. The cochlea is divided into two sections by the **basilar membrane [to glossary]** which runs down the middle of the cochlea. The basilar membrane attaches to the wall of the cochlea right beneath the oval window. This end of the cochlea is called the **base [to glossary]**. The other end which reaches almost to the tip of the cochlea is called the **apex [to glossary]**. The space at the end of the cochlea where there is no basilar membrane is called the **helicotrema [to glossary]**. Along the basilar membrane are those structures that perform transduction of the auditory stimulus. It is therefore necessary to take a closer look at the basilar membrane and the structures on it.

Figure 10.x shows a cross section of the cochlea. Upon the basilar membrane is a set of structure that are collectively referred to as the **Organ of Corti [to glossary]**. Covering the Organ of Corti is another membrane, the **tectorial membrane [to glossary]**. The tectorial membrane covers the receptor cells of the auditory system. These cells are the **hair cells [to glossary]** and they are called hair cells because of the cilia that protrude from the top of each cell. There are two sets of hair cells, the **inner hair cells [to glossary]** and the **outer hair cells [to glossary]**. The outer hair cells are given that name because they are near the open end of the tectorial membrane. The inner hair cells are nearer the attached end of the tectorial membrane. Figure 10.x shows a close up photograph of those cells. There are three rows of outer hair cells and one row of the inner hair cells. Each cell has many of the cilia and as can be seen from the photograph, they cilia are arranged in a V shape. This shape to the cilia will become important. The longest of the cilia is the cilia at the peak of the V and it is called the **kinocilium [to glossary]**.

What does the Sound Stimulus do to the Inner Ear?

It is now time to bring the stimulus from the motion of the stapes up to the point where auditory transduction takes place. In other words, at this point the stapes is pressing against the oval window, but what does that do to the inner ear. The first clear answer was found by Georg von Bekesy (XXXX) who attached a piston to the oval window and started it moving in and out like the stapes would and he observed what happened in the inner ear directly. Bring up **Interactive Illustration 10.x, The Cochlea** and press the **Start Sound** button to simulate what von Bekesy observed. This button will start the stapes pressing against the oval window and will model the behavior of the inner ear. When the stapes presses in against the oval window it increases the pressure in the fluid in the cochlea in that area. When the stapes pulls out, it decreases the pressure in the fluid in the cochlea in that area. These pressure variations in the fluid are transmitted down the cochlea and also onto the basilar membrane. When the pressure of the fluid increases, it pushes down the basilar membrane, and when the pressure of the fluid decreases, it pulls up the basilar membrane. Thus, the basilar membrane starts moving in a wave like fashion, ultimately tied to the sound wave stimulus. This wave moves along the length of the basilar membrane in what is called the **traveling wave [to glossary]**. What this means at the hair cell is that they are moving up and down on the basilar membrane and the cilia are flowing through the water in both directions.

At this point, a small detour will be taken to step back and look at the complexity of the ear. To review, there are three main portions of the ear: outer, middle, and now inner ear. To this point, the receptor has yet to be encountered. One question that might be asked is why all of this complexity. In engineering, an important principle is the KISS principle. Keep It Simple Stupid. The more complex the system the more likely something will go wrong. So making a system more complex should only be done for the best of reasons. Is there a reason for the complexity of the ear? The relative density of liquids compared to air suggests an answer to this question. If the air pushed directly against the oval window it might not move the fluid very much. So the larger size of the oval window to the bottom of the stapes and the motion of the ossicles themselves act as an amplifier to compensate for the difference in density between the air and liquid of the inner ear.

Review of the Auditory Stimulus in the Ear

Open **Interactive Illustration 10x: The Ear** to review the path of the auditory stimulus from the outer ear to the inner ear. This image shows both the middle and inner ear connected to the end of the external auditory canal. Pressing the **Start Sound** button will start a sound stimulus entering the outer ear and show how this stimulus triggers the events that have been discussed in the middle and inner ear. To review the middle and inner ear alone, you can select them using the menu at the top of the screen. The sound stops each time you change the screen so you will have to restart it. You can also adjust the **Frequency** and **Amplitude** of the sound stimulus as well. Examining these variables will anticipate some questions to be addressed later.

Auditory Transduction

[I WILL NEED SOME ANIMATIONS FOR THIS AS WELL. I MIGHT BE ABLE TO DESIGN SOME]

It is at this point that auditory transduction takes place as a result of some fascinating features of our anatomy. Figure 10.x shows a representation of the cilia of one hair cell. The kinocilium is the tallest of the cilia. Each of the cilia is connected to each other by a protein bridge called a tip link. The tip link covers ion channels and these ion channels will be crucial in transduction.

Open **Interactive Illustration 10.x, The Hair Cell and Auditory Transduction** to see an animation that will illustrate how the motion of the basilar membrane is translated into auditory transduction. When the traveling wave reaches a particular portion of the basilar membrane, those hair cells move up and down with the basilar membrane. Flowing through the fluid in inner ear, the cilia are pushed in the opposite direction of the motion of the hair cell. This flowing motion is called a **shearing force [to glossary]**. Press the **Start Sound** button to start the movement of the haircell in response to the traveling wave along the basilar membrane (not shown). The direction of the shearing force on the cilia is indicated on the diagram. The hair cells are not all the same length. The kinocilium is the longest and the hair cells get shorter as they go away from the kinocilium. Since it is the tallest, the kinocilum catches the most of the fluid and is pushed the hardest. It is similar to the tallest trees in the woods catching the most of a breeze. So, when the hair cell shown in the interactive illustration moves to the left the kinocilium is pushed hardest to the right. As a result of this motion, the cilia pull apart pulling the tip links tight. When the tip links are pulled tight, they uncovers ion channels allowing the hair cell to depolarize which causes it to release neurotransmitters to the excite the neurons that it connects to which make up the auditory nerve. When the hair cell moves to the right, the cilia are pushed together. It is only the motion in the one direction that stimulates the hair cell and it is this motion that causes transduction. It is worth noting, because it will be important later, that the ion channels open once each cycle of the sound wave.

The motion of the hair cell and cilia are all slowed down to help you see the motion, but it is a complex motion and the details can be hard to follow. To see the motion in different lings you can remove and restore the vertical component of the cell motion but clicking on the **Show Vertical Motion** checkbox at the bottom of the screen. You can also remove and restore the tilting of the cell as it rides the basilar membrane wit the **Show Tilt Motion** checkbox. You can also slow down or speed up the motion of the hair cell with the **Frequency** slider at the bottom of the screen.

[I NEED SOME COMPARATIVE INFORMATION HERE – I HAVE SOME FROM LAST YEAR FIND IT]

[SAVE THE BELOW DISCUSSION FOR AFTER THE DISCUSSION OF THE EAR]

Ohm's Acoustical Law. The ambiguity about whether a sound frequencies are hears separately or only contribute to the color experience of the sound in the discussion of timbre earlier in this chapter leads to one of the more interesting abilities in hearing that needs to be explained. While in many cases we can hear a complex sound of many frequencies as a single tone, we can also break that tone down and often hear the individual frequencies. Remember from the discussion of the ear, there is only one place where the sound stimulus enters our auditory system. All sounds that we hear at any one time all must ultimately cause the stapes to which causes, in essence, a single stimulus to enter the ear at any one time. This fact is very different from the visual system were two stimuli can stimulate different locations on the retina due to the focusing of the cornea and lens.

Yet, in the ear we often separate different stimuli easily and this separation can even occur when there are different frequencies that make up a complex tone. Open **Interactive Illustration 10.x, Ohm's Acoustical Law** **[link to media]** to experience this ability. Looking at the graph of the sound wave, you are seeing a stimulus that is made up of 10 component frequencies that are whole multiples of the fundamental of 200 Hz. Below the graph of the sound wave are controls to control each component of the sound stimulus. The components are indicated by the numbers **0** through **9**. The **0** refers to the fundamental. Each higher number is the next higher frequency component of the sound. The **1** refers to the second component and has a frequency of 400 Hz at the start of the illustration. For each, component you can adjust the frequency using the **Frequency** slider right below the number or remove and add the component using the **Cancel** checkboxes below the slider. The **Reset** button will return the illustration to its starting settings.

Click the **Play** button to start the tone. When the tone first comes on most of us will hear a single tone and not be able to discriminate the different frequencies, though experience musicians can. That is, generally, we hear the synthetically. For the rest of us, a little help can give us the ability to hear the separate frequencies or hear the tone analytically. Then click the **Cancel** checkbox for the first component and remove this component by cancellation with copy of the tone that is 180 degrees out of phase. Then click the **Cancel** checkbox again to restore the component. The removal and addition of the component should allow you to focus on this component and hear is separately from the rest of the tone. Repeating this task on the other components than the fundamental will allow you to hear each component in turn. The fundamental is already generally heard because it is the tone that determines the pitch of the complex so it does not need the help to be heard separately. You can also change the frequency of each component to see if changing frequency, like changing loudness, enables you to pick out the tone. If it does, are there any conclusions about what you need to hear a component from the background?

Fourier Analysis in Audition. Ohm's Acoustical Law suggests that the auditory system has an ability to take a complex waveform and break it down into its component frequencies. This ability is like Fourier Analysis discussed for vision. Open **Interactive Illustration 10.x, Spectrum Analyzer** **[link to media]** to see how the results of a Fourier Analysis will be represented for sound stimuli. There are two graphs in this figure. The top graph is the standard graph of the sound wave that has been used several times already. The bottom graph is a graph that you will see often from now on. It is a graph of the result of a Fourier Analysis of the sound wave shown in the top graph. The x axis is the frequency and the y axis is the relative amplitude of the different components. On the left side of the screen are sliders to control a few frequencies. When the figure first comes up, there is only one frequency present and that is shown in the spectrum graph by a single peak. Using the second slider (labeled **1**) causes a second peak to appear in the spectrum graph. As each new frequency is added a new peak appears on the spectrum analyzer graph. Thus, while the graph of the sound wave is a continuous variation in sound intensity, the spectrum graph shows distinct peaks for each frequency showing that the Fourier Analysis that drives the spectrum graph has separated each frequency, much as we can as shown my Ohm's Acoustical Law.

To make the similarity between Fourier Analysis discussed for the visual system and what is being discussed here more precise, open **Interactive Illustration 10.x, Fourier Analysis in Audition** **[link to media]**. In this figure, you have the same graph of the sound wave below with the same spectrum graph below. There is a **Wave Shape** menu at the top left corner of the screen. Currently it is on **Sine**. This waveform and spectrum should look familiar from the last figure. Select **Square** and a close approximation of a square wave will be shown (the computer cannot exactly reproduce a square wave sound). You will see on the spectrum graph a series of peaks at regular intervals across the spectrum graph. Each peak will be the odd multiple of the fundamental which is the largest peak on the spectrum graph. The greatest intensity is for the fundamental and each higher harmonic has a lower intensity. This graph should look very similar to the graphs for the square wave shown in when the visual system was discussed. You can also pick **Triangle** for a waveform and compare it to the triangle wave when the visual system was discussed.

There is one more sound stimulus that can be used in this figure. This sound stimulus is called **Noise**. This is white noise and is a random and constantly varying collection of frequencies in the stimulus. To experience this sound, which will be used several times in the examples to be used later in the chapters on audition, press the **Play** button. The stimulus sounds like static and static is a noisy stimulus. If you press the **Reset** button under the sound stimulus graph you will see that the waveform is constantly changing. That would not be true of the other stimuli. If you press the **Retrigger** button under the spectrum graph, while the sound is playing, it will trigger a new Fourier Analysis on the current sound stimulus. Each time a new Fourier Analysis is done, a new patter shows up on the spectrum graph. Give the graph a little time to update. The Fourier Analysis is a complex mathematical equation to compute and it takes a bit of time to complete the operation, which is one reason that the graph does not automatically update. These changes in the spectrum graph show that the sound stimulus is a collection of all of the sound frequencies possible and that the intensities of each of these frequencies are constantly changing.

Auditory Pathways

The sound stimulus follows an interesting path from the ear to the brain. The hair cells along the basilar membrane connect to the neurons that make up the auditory nerve which is part of the VIII Cranial nerve. Interestingly, though there are many more outer hair cells than inner hair cells, about 95% of the nerves from that make up the auditory nerve start from the inner hair cells (Green, XXXX). From this observation, it is clear that the inner hair cells are far more important to our conscious experience of hearing.

Figure 10.x shows the auditory pathways. The auditory nerve enters the brain at the level of the medulla and then synapse in the **cochlear nucleus [to glossary]**. At that point the sound path splits. Some of the sound pathways decussate or cross to the contralateral or other side of the brain and some of the sound pathway stays on the ipsilateral or same side of the brain. Most of the fibers cross to the other side (about 2/3rds). In both cases, the target of the auditory fibers is another structure in the medulla, the **superior olivary complex [to glossary]** where the pathway runs into another synapse. These structures of the medulla are not merely stop ways along the path to the cortex. These structures perform many functions, but one that will be important to our later discussion is that these structures seem to play some role in the processing of sound necessary to determine the perceived direction of the sound.

From the superior olivary complex, the auditory information on both sides of the brain travel to the **inferior colliculus [to glossary]** for the next synapse. This structure is located in the midbrain, a very small region of the brain. The inferior colliculus seems to play an important role in helping us reflexively orient to a sound. So if there is a large sound and we turn towards that sound, the inferior colliculus has played an important role in that process (REF XXXX). From the inferior colliculus, the auditory pathways travel to the **medial geniculate nucleus [to glossary]** in the thalamus for the last synapse before reaching the cortex. The name of the medial geniculate nucleus might ring a bell. The lateral geniculate nucleus is nearby, mostly just a little bit farther to the side of the brain, and is main part of the thalamus where visual information travels to on its way to the cortex. [I NEED SOME SENSE OF THE ROLE OF THE MGN]

From the medial geniculate nucleus, the auditory information travels to the **superior temporal gyrus [to glossary]** which is on the side of the brain. Figure 10.x, shows the auditory cortex from the side of the brain. The auditory cortex is in the temporal lobe which is the large thumb-like protrusion on the side of the brain that is largely separate from the rest of the brain. A gyrus is one of the bumps on the side of the brain and the superior temporal gyrus is the top gurus. So the auditory cortex is on the top part of the temporal lobe. The auditory cortex performs many functions related to our hearing. One point that is worth describing here is the finding that the auditory cortex is organized. One way that the auditory cortex is organized is tonotopically, that is, the auditory cortex has a **tonotopic map [to glossary]**. A tonotopic map is where the layout of the auditory cortex is arranged by the range of pitches. At one end of the auditory cortex will be neurons sensitive to low pitches, or frequencies, and at the other end of the auditory cortex, will be neurons sensitive to high pitches. In between, the neurons are arranged in such away as you can layout the pitches in order across the surface of the auditory cortex.

Fundamental Dimensions of Sound

The perception of sound is a complex and rich experience. To start this discussion, a couple of topics will be discussed that are very basic and tied closely to the fundamental measures of the sound stimulus, intensity and frequency. Regardless of the complexity of the sound, all sound experiences are impacted by these dimensions of auditory experience and to some extent these dimensions must be represented or encoded in the brain so that they can be experienced.

Encoding Intensity

Intensity is most closely related to intensity so this section will discuss how loudness relates to intensity. Generally loudness like brightness does not increase as fast as intensity. The relationship of intensity to loudness is well described by Steven's Law. If you need a review of Steven's Law open **Interactive Illustration 10.x, Review of Steven's Law** **[link to media]**. Recall the Law:

$$S = cI^b \qquad (3)$$

S stands for the strength of the sensation in this case the how loud the sound it. I is the intensity of the sound and b, the exponent, is the more important part of the function. If b is less than 1, like the default in **Interactive Illustration 10.x, Review Steven's Law** then the as intensity increase, the sensory strength goes up ever more slowly. The decibel scale takes this feature of auditory perception largely into account by the use of the logarithms.

Sound Intensity doubles every ~6 dB but if you play **Interactive Illustration 10.x, The Decibel Scale** **[link to media]** you will find that sound loudness does not double when the amplitude of the sound doubles. When you open this interactive illustration, you will see graph to show sound waves at the top and a bar graph on the bottom. In this interactive illustration, you will press the **Start Series** button to initiate a series of short tones where each successive tone is less intense by the number of decibels shown in the **Decibel Step Size** menu. The initial **Decibel Step Size** is **5** dB. The **Intensity Sequence Graph** bar graph shows the intensity of the whole sequence of sounds that will be presented in two ways. The cyan bars show the **Relative Amplitude** of the sound waves. The blue bars show the difference in the relative intensity in decibel steps. Each **decibel Values** bar is an equal size step down showing that each sound will be an equal number of decibels down. When you press the **Start Series** button, the current sound will be shown as a green wave form in the **Soundwave Graph**. To help you keep track of where you are in the sequence, the bars in the **Intensity Sequence Graph** for the sound that is playing will also turn green with the **decibel Values** bar turning a darker green. After the first tone plays, the current tone will still be shown in green, but the immediate past tone will be shown on the **Soundwave Graph** and the **Intensity Sequence Graph** in red to help you compare those two sounds.

Press the **Start Series** button and listen to the sequence of tones and follow the sequence on the two graphs. The first think you will notice is that each sound seems to be an equal amount softer than the one sound preceding it. This pattern follows the decibel scale and not the amplitude scale directly. With the amplitudes, the difference between each successive pair of tones gets smaller but their proportional difference, captured in the decibel scale, stays the same. This finding that proportional differences are important to hearing differences in sound intensity and not absolute differences in intensity results from the exponent in the Steven's Power Law for audition being less than one.

Now try a variation on this sequence. Use the **Decibel Step Size** menu to select the value of **6**. A difference of 6 dB is where the amplitude of the more intense sound is very close to twice the amplitude of the less intense sound. If you look at the cyan bars of the **Intensity Sequence Graph** you will see that second cyan bar is half the height of the first cyan bar. You can use a ruler and measure it. Now, play this sequence and see if each successive tone is one half as loud as its predecessor?

After you have completed this sequence, select other **Decibel Step Size** values and try them out to get a feeling for what different size decibel steps sound like. You can also select different frequency values to compare the decibel scale for different ranges of frequencies.

To directly compare intensity and loudness try **Experiment 10.x, Find a Sound Half as Loud** **[link to media]**. This is a very simple experiment that should not take very much time. The experiment is designed to allow you to play one sound as the standard and then adjust the intensity of the comparison until it sounds half as loud to you as the standard. You can play the standard by clicking on the **Play Standard** button. You can play the comparison by clicking on the **Play Comparison** button. You can adjust the intensity of the comparison between trials by using the **Comparison Amplitude** slider below the two play buttons. You can also play the two tones in succession with the standard going first by pressing the **Play Pair** button. Adjust the intensity of the comparison until it sounds half as loud to you as the standard. When you are satisfied with your choice, press the **Comparison is Half** button. When this button is pressed the two amplitudes are graphed on the **Relative Intensities** bar graph and your data are summarized in the **Output Table**. The **STD Gain** value is the relative intensity of the standard and is always **1.0**. The **CMP Gain** value is the relative intensity of the comparison and is determined by the **Comparison Amplitude** slider. The **dB Diff** value is the difference between the two intensities converted into decibel units. The bottom line of the **Output Table** are the values you should find if your find the comparison to be half as loud when its intensity is actually half that of the standard. The **CMP Gain** is **0.5** and the **dB Diff** is **6.02**. Most people find that their **CMP Gain** is less than **0.5** and their **dB Diff** is **6.02** indicating that physical differences are larger than perceptual differences, which is what is expected from a Steven's Law equation with a b less than 1.0.

Auditory Thresholds. To understand our perception of loudness it is important to first have some understanding of auditory thresholds. Open **Experiment 10.x, Frequency Response of the Ear** **[link to media]**. This experiment will allow you to perform a quick experiment to measure your auditory thresholds. By doing this experiment, you will understand a little more how such data are collected and collect some data to understand some of the character of auditory thresholds. To do this experiment at its best you need a quite place and to turn off other sound sources such as any music or movies or television. It would also be wise to not use speakers that are on a notebook computer. These speakers, while getting better, do not have a sufficient frequency response for this experiment. So either use headphones or use speakers that are used with desktop computers. Those speakers are usually quite adequate to these purposes.

When you first open the experiment, there is a brief calibration screen. You will need to play the tone by pressing the **Play** button at the top of the screen. Use the volume control on your computer to adjust that sound so that you can just hear it. This adjustment assures that the range of sounds that will be played will work to help determine your auditory threshold. When you have completed this adjustment, press the **Stop** button to stop the sound and the **Done** button to proceed to the experiment.

When the experiment screen comes up, there is a sound wave graph like used before at the top of the screen and across the bottom is a graph that will display your data. The x-axis of the data graph is frequency (in logarithmic steps) and the y-axis is the intensity of sound in dB with the comparison dB being the loudest sound used in the experiment which is why the top of the y-axis is 0 and the rest of the y-axis is negative. There are two controls on the left side of the screen. The first control is the **Frequency** menu which allows you to pick a frequency to test your threshold. The next control is the **Start Series** button to begin a descending stair case of sounds. When you press this button, a sound will play at the frequency selected in the **Frequency** menu and at the loudest intensity used in the experiment. The sound wave will be plotted on the sound graph so that you can know that the sound is playing. This sound will play for a short period of time. There will be a brief break and then the next sound will play and it will be 5 dB softer than the first sound. There will 10 of these steps in the experiment. Your task will be count how many of these steps you can hear. Go ahead and run the first staircase at 125 Hz.

When the staircase is done, several more controls appear. Right below the **Frequency** menu and the **Start Series** button are controls to allow you to enter the number of steps of the staircase that you have heard. Fist is the **Steps Heard** menu. Use this menu to enter the number of steps of the staircase that you have heard. When you have the proper number in your menu, press the **Data Entered** button right below the **Steps Heard** staircase. If you need to you can press the **Start Series** button again to replay the series. When you have pressed the **Data Entered** button, the data that you have entered will be displayed on the data graph. If you have heard 1 of the 10 tone and enter a 1 for the number of steps you have hear, a -5 will be plotted on the graph, since each step is 5 dB below the loudest. If you heard 3 steps, the value entered on the graph will be -15. After entering data, you can use the **Frequency** menu to select a new frequency or run the same staircase again. If you run the same stair case and enter new data, the graph will plot the arithmetic mean of all of the times you have entered data for that stair case.

At the bottom left side of the screen are two more buttons, **Clear Data** and **Clear Frequency**. The **Clear Data** button will clear all of the data from the frequency that has just been run. The **Clear Data** will clear all of the data that has been collected and return you to the situation you had at the beginning of the experiment. The **Clear Frequency** button is to clear data for an individual frequency if you enter the wrong data by accident. The **Clear Data** button can be used if you need to start over for some reason.

To test your thresholds at different frequencies run at least one staircase at each of the 8 frequencies that can be tested (**125**, **250**, **500**, **1000**, **2000**, **4000**, **8000**, **10000** Hz). Running more than one staircase will more likely give you regular data and it is recommended that you run at least two staircases. When you have run at least one staircase for each frequency, return to the book.

You data should look something like Figure 10.x. Generally, people are most sensitive to frequencies around 1000 Hz to 4000 Hz. People are less sensitive to higher and lower frequencies. Part of the reason for the greater sensitivity in this range of frequencies is that the shape of the pinna and external auditory canal allow these frequencies to pass through more easily. The pinna and external auditory canal are said to have a resonance frequency in this range.

Equal Loudness Contours. Fletcher and Munson (1933) did an important experiment about the relationship between loudness, intensity and frequency of the tone. It is worth giving an outline of how this experiment was run before describing its results and the implications of those findings. In brief, a participant was presented a tone of 1000 Hz at a certain intensity, say 10 dB. This tone would be the standard for this condition. The participant's task was then to adjust the intensity of tone of a different frequency, for example 500 Hz, until the two tones sounded equally loud. This tone is the comparison. The frequency of the comparison was changed and the process repeated. Eventually, the standard was compared with a set of frequencies all across the audible range. The data from this condition of the experiment could be plotted on a graph like in Figure 10.x. The dotted line represents the auditory thresholds you have seen before. The solid line labeled 10 is what is called the equal-loudness contour. All the points along this line sound equally loud to the participants. For example a tone of ~25 Hz at an in intensity of 70 dB sounded equally loud to a 1000 Hz tone at 10 dB and also a ~7500 Hz tone at 20 dB.

This experiment by itself would require a lot of time on the part of the participant, but so far only one condition has been described. The experiment was repeated but for different intensities of the 1000 Hz standard tone. Figurer 10.x shows the complete results of the Fletcher and Munson (1933) experiment. Each line on the figure is another equal-loudness contour. Now, the contours are not precisely parallel to each other. The lines are closer together at the low frequency end and the high frequency end.

Open **Experiment 10.x, Equal Loudness Contours** **[link to media]**. In this experiment you will get to do a very small portion of a Fletcher and Munson type of experiment. This experiment is best run with either headphones or in a very quiet room. When the experiment comes up a calibration window is opened. The tone needs to be audible but soft. This is the softest version of the standard tone that will be played. When the calibration tone has been adjusted press the **Done** button to proceed to the experiment. The screen is very similar to **Experiment 10.x, Find a Sound Half as Loud**. On the upper left hand corner of the screen you can adjust the condition to be tested. The top center portion of the screen is the

Soundwave Graph to show the sounds being played. Below the **Soundwave Graph** is a graph and table to show your results as you generate them.

The standard tone will be 1000 Hz like in the Fletcher and Munson Experiment. With the **Frequency** menu, the frequency of the comparison can be selected. The intensity of the standard frequency can be selected by using the **Standard Intensity** menu just below the **Frequency** menu. There are four values than can be selected and each is given in terms of the dB of the tone relative to the lowest intensity used, which is why there is a **0** value on the menu.

In this experiment, the standard will be played at the set intensity followed by the comparison at an intensity decided upon by the experiment. This pair of tones will be repeated with the comparison being played at every decreasing intensities. Your task is to count in how many of these pairs the comparison sounds louder than the standard. There are 9 pairs of standard followed by comparison tones in a series. Press the **Play Series** to start the first condition. If you have not changed any of the settings, the **Frequency** of the comparison will be **60** Hz and the standard intensity will be **0** dB. As the series plays, count the number of times the comparison sounds louder than the standard.

After the series has played, the data collection objects are turned on. Enter the number of steps louder in the **Number of Steps Comparison is Louder than Standard** menu and then press the **Enter Data** button. When this button is pressed the **Relative Intensities** graph will be updated to indicate your data entry with the data being converted to decibels. The value will also show up on the **Output Table**. There are already data entered for the 1000 Hz frequency as this is the standard and set up by the experiment. Each intensity of the standard is given a different line in the **Relative Intensities** graph and a different row in the **Output Table**. If you When you have entered the data, you can either repeat the condition or chose a new condition by varying either the **Frequency** of the comparison or the **Standard Intensity**. If you repeat a condition, when you enter the data, the value plotted on the graph and table will be the mean of all of the runs of that precise trial. You can clear all of the data and start over at any time by pressing the **Clear All Data** button. After you have entered data, if you figure you have made a mistake you can remove the last value entered by pressing the **Clear Last Value Entered** button but only before you change any other values or run a new series. Go through all of the conditions of the experiment and fill out the graph and table and then return to the text.

Looking at the graph, most people find that the **0.00** dB data looks like a V with the **60** Hz and **10,000** Hz tones having to be much more intense to sound equally loud as the **1000** Hz tone. The V might be a little steeper on the **60** Hz side reflecting that auditory sensitivity falls off faster on the low frequency end of the audible range of frequencies. This finding is not surprising. It follows from both the auditory threshold and the Fletcher and Munson data. First, look at the graph and see if you see any resemblance between your data and the data found by Fletcher and Munson shown in Figure 10.x. Fletcher and Munson's data is replicated if the V for the **40.00** dB data is flatter than the **0.00** dB data. This replication can also be verified if in the table by examining the differences between the intensity of the comparisons and their matching standard, which is on the same row of the table. The differences at **40.00** dB should be smaller than at **0.00** dB.

[FOR NOW ONLY A DEMONSTRATION VERSION OF THE EXPERIMENT IS DONE. A MORE COMPLETE VERSION NEEDS TO BE ADDED]

There are some important implications of these findings. The most important implication has to do with replaying a recording. If music is recorded at one set of intensities and then played at a dramatically different intensity, the sound can be quite changed. An artificial example will suffice. Say some music was recorded that had a 50 Hz tone and a 1000 Hz tone that were both equally loud and the 1000 Hz tone is 100 dB. To be equally loud the 50 Hz tone is 110 dB. Now, you play back that music at a much lower intensity or volume so that you can use the music as background music. Say the 1000 Hz tone is 30 dB. To be equally loud the 50 Hz tone would have to be about 50 dB but it is only 40 dB. Not only is this tone not equally loud but it is below the auditory threshold so it will probably not be heard at all.

Some stereos have a quick fix for this issue. On some older and not the most expensive stereos there is a loudness button. This loudness button boosts the intensity of both the lower and upper frequencies to help balance the loudness of these frequencies when listening music at a low volume. More sophisticated stereos have equalizers that allow the listener to adjust many different frequency ranges to adjust intensity for many circumstances.

Encoding Frequency

The frequency of the sound is most closely related to the psychological experience of pitch. Pitch is that quality of sound that we report as low or high. The layout of a piano keyboard is a good illustration of pitch. So, the ability process frequency is the foundation of much of pitch perception which will be discussed more fully in the next chapter.

Frequency Discrimination. The ability to discriminate between two frequencies is described very nicely by Weber's Law. Discrimination, if you recall, is also the difference threshold or Just Noticeable Difference (JND). JNDs are proportional to the base stimulus that is being discriminated against when JNDs follow Weber's Law. In practical terms, this means that it easier to discriminate between two low frequencies than two high frequencies if they are equally separated. The Weber's Constant for frequency discrimination is 0.004. Thus, the average listener can discriminate between 500 and 502 Hz reliably. However, if the base stimulus is 5000 Hz, to reach JND, the comparison stimulus must be 5020 Hz. Though it is seems that discrimination may be dependent upon many factors leading to some confusion in the literature Green (1976).

Open **Experiment 10.x, Frequency Discrimination** **[link to media]**. In this experiment, you will measure your own frequency discrimination and you can compare it to the standard value given above. When you first open the experiment you will be presented with a screen to allow you to set up the stimulus parameters of the experiment. You can set parameters such as the **Number of Conditions to Test,** the frequency of the standard stimulus for each of these conditions, the **Duration of the Stimuli**, the **Upper Limit of the Staircase** or the greatest difference in frequency to be tested, the **Interstimulus Interval** or time between the standard and comparison, and the **Intertrial Interval** or the time after the completion of trial before the next trial begins. These values can be left at the defaults for this time through. Next, the method screen that you have seen before will be shown. You will be running a staircase method if you leave the settings at the default. One default value is worth mentioning. The **Number of Levels to Test** is set to **15** different stimulus levels because of the running of several different conditions in this experiment. The next screen will be the calibration screen. In this case, you just need to set the intensity to a comfortable level. There is no need for precision.

After the calibration screen, the experiment begins. The standard plays first followed by the comparison. Their waveforms will be displayed on the **Soundwave Graph** at the top of the screen. You do not need to watch them carefully but they may help you to know when each sound is playing. As before, the current sound is drawn in green and the past sound is in red. So when the standard is playing, there will be one waveform and it will be in green. When the comparison is playing, the comparison will be in green and the standard in red. When the trial is done your task is to answer the question, "**Did the second tone have a higher pitch than the first?**" If it did press the **Yes** button or **z** key. If it did not, press the **No** button or **/** key. When you have responded the next trial will begin after the intertribal interval.

When you have complete one condition, the data from that condition will be shown in the **Output** portion of the screen. The **Frequency Differences** graph will show a plot of the size of the JND for this condition as the frequency of the discrimination threshold minus the frequency of the standard. In the **Output Table** the frequency of the threshold (**Thresh**), the frequency difference (**Diff**), and the size of the Weber fraction (**Weber**) are all displayed. Run the experiment and then return to the text.

Based on Green (1976) you might see some variability in your responses. However, the frequency difference should generally be greater for the high frequencies of the standard than the lower. If Weber's law holds tightly you for you, you should see a straight line that ascends from left to right, that is, it has a positive slope. The Weber fractions in the table should all be close to each other. You might get many results near the 0.004 value given above for the Weber Fraction. Values that are double or half this standard fraction are still within range given the variability of this task. Still, you should see that on an absolute scale discrimination of frequencies gets harder as the frequency of the standard gets higher.

These findings do not indicate how the brain is able to hear different frequencies as different psychological experiences. The question at this point is how does the ear capture some information about the frequency of the sound so that our brain can use it. Like color vision it seems that there are two possible theories that are involved in this process. Each theory will be explored from the angle of the questions asked and observations made that led to those conclusions.

Place Theory. This theory starts out by observing some interesting features of the traveling wave that has been glossed over so far. Reopen **Interactive Illustration 10.x, The Ear** to review the traveling wave. Press the **Start Sound** button at the bottom of the screen. When the sound is moving, the molecules of the air change density which moves the tympanic membrane. The movement of the tympanic membrane is translated to the ossicles of the middle ear. The stapes pushes against the oval window causing the traveling wave to move along the basilar membrane. Use the **Frequency of Sound** slider to adjust the frequency of the incoming sound. What happens to the traveling wave as the frequency gets higher?

Without much description, it is obvious from this observation that the traveling wave is changed by the frequency of the incoming sound. Before going forward with this discussion, it would be useful to try to understand how this is possible. It goes back to a concept that has been discussed earlier in other contexts.

To get an answer to this question, an examination of the structure of the basilar membrane will be helpful. The basilar membrane is very interestingly constructed. It is not identical all along its length. Near the base, the basilar membrane is narrower and stiffer. Near the apex the basilar membrane is wider and less stiff. These differences are important and play a role in the ability of the ear to discriminate frequency. These differences in the structure of the basilar membrane affect the resonance of each position of the basilar membrane. Recall, that for any matter, it has a preferred frequency and this is the resonance of that matter. Generally, the shorter the length of an object or the stiffer it is, and, thus, it tends to resonate at a higher frequency. Think of a guitar. When the player plucks a string is vibrates most strongly at one frequency which is the pitch that is produced by that string. The play can change that resonant frequency by putting a finger on one of the frets of the guitar. This changes the length of the string, making it shorter. Also, the frequency gets higher.

[GET A SHORT VIDEO TAPE ILLUSTRATING THIS?]

Now, return to the basilar membrane and how it is constructed could play a role in the encoding of frequency, which is the main question of the moment. To answer this question, it would be nice to be able to plot how the traveling wave affects each point of the basilar membrane for all sorts of different frequencies. This is precisely what von Békésy did in some ground breaking studies to examine audition.

Returning to the basilar membrane, the different properties of the basilar membrane along its length means that different positions will have different resonant frequencies. Each point will vibrate most easily to different frequencies. von Békésy (1957, 1960) was the first to witness the impact of the construction of the basilar membrane on how sound is processed in the inner ear. He extracted cochlea and opened them. To simulate sound stimuli, he attached a very fine piston to the oval window which he could move similarly to a sound of a single frequency.

Interactive Illustration 10.x, Traveling Waves and the Basilar Membrane will allow a simulation of this experiment. Open the illustration and work through the activity as described here and see if a hypothesis about how the motion of the traveling way can lead to our hearing different frequencies. First, let me describe what you see on the screen. Near the top of the screen is a horizontal white line. This represents the **Basilar Membrane**. The rest of the cochlea and hair cells are not drawn to help show the motion of the basilar membrane. The **Base** and **Apex** of the basilar membrane are labeled on the screen. Below the basilar membrane is a graph that will plot how far the traveling wave moves the basilar membrane from its resting position for each position on the basilar membrane. This distance, plotted on the y axis, will be referred to as the height of the traveling wave. The x axis is the representation of the basilar membrane with position 0 being the base and the other end of the graph being the apex.

In this experiment, different frequencies of sounds will be applied to the oval window and a graph of the height of the traveling wave as it travels along the basilar membrane will be generated. On this graph you can plot the height of the traveling wave for up to three frequencies simultaneously. Each frequency will be plotted in a different color (red, yellow and green). To ease the use of the illustration, some preset values for the frequencies of each curve on the graph have been created, though you can choose any value of frequency for any curve by using the **Frequency of Sound** slider to select the frequency of the sound and clicking on the desired **Curve To Plo**t checkbox next to the **Frequency of Sound** slider.

To start the experiment, press the **500 Hz** button to the right of the graph of the hight of the traveling wave. This button will simulate the basilar membrane with a simulated 500 Hz tone and prepares to plot the height of the traveling wave for the first (red) line on the graph, clearing any previous data. If you look at the legend for the graph, the red line will be labeled with **Freq: 500**. When the basilar membrane moves, a dashed line is drawn at the point where the basilar membrane rests when there is no sound. This dashed line can serve as a reference. The height of the basilar membrane is the distance it travels from that position. To add a data point to the plot, move the mouse pointer to over the region where the basilar membrane is being drawn. A vertical blue line will be drawn to indicate the position of the mouse along the basilar membrane. A dashed blue line will be drawn on the graph to show the same relative position on the graph. Clicking while the mouse is over the region of the screen that draws the basilar membrane will plot a point on the graph to show the farthest the basilar membrane moves from its resting point at that position. This is the height of the basilar membrane at that point.

The height is automatically determined because the **Auto Peak** option has been selected in the **Auto Peak** menu. Unselecting this option by clicking on **Off** in the menu will require you to manually find the farthest distance from the basilar membrane with the mouse. To aid this process, the traveling wave can be stopped in its track with the **Hold Sound** button and a horizontal guide is drawn.

For now, leave the **Auto Peak** option selected to ease the determination of the curve. To fill out the curve, move the mouse up and down along the basilar membrane and keep clicking. When you think you have a good representation of the curve, click on the **1500 Hz** button and repeat the process. This button will have the data plotted on a new line on the graph so that you can compare the effects of these two frequencies directly. Clicking on the **5000 Hz** button will allow you to plot a third set of data on a third line on the graph.

If you want to choose your own frequencies, use the **Frequency of Sound** slide as mention above. It is still possible to plot up to three set of data if whenever you change frequencies, you select a different **Curve to Plot** checkbox (**Curve 1**, **Curve 2**, or **Curve 3**). If you need to you can clear a set of data use the **Reset** button which will reset the data from the selected **Curve to Plot** checkbox.

Using either the present frequencies or three of your own (one low, one moderate, and one high frequency) plot the height of the traveling wave for the three frequencies and compare them. In particular note where the traveling wave for each frequency moves the farthest from the basilar membrane (that will be the peak of the associated curve on the graph). Also note how sharp and well defined the peak of the curve is.

These results will simulate the findings of von Békésy (1960). Von Békésy found that the traveling wave peaks depends upon the frequency of the incoming stimulus. Recall that shorter, stiffer objects tend to have higher resonant frequencies. In high frequency stimulus (5000 Hz) example, the traveling wave climbs quickly to its peak and then falls off quickly. At the frequency of the sound gets lower (1500 Hz to 500 Hz), the peak of the traveling wave height curves moves further down the basilar membrane towards the apex.

These results lead to a conclusion in two parts. First, different frequencies of sounds will cause the traveling wave to move the most at different locations along the basilar membrane. Second, the positions of these peaks are in a nice ordered relationship: the lower the frequency, the closer the peak to the apex; the higher the frequency the closer the peak to the base.

So what does it mean that different frequencies will cause the traveling wave to peak at different places? Recall that the hair cells are arranged all along the basilar membrane. The results if the interactive illustration leads to the conclusion that different hair cells will be stimulated by different frequencies. If different hair cells are stimulating different neurons of the auditory nerve then different frequencies ultimately stimulate different neurons going to the brain. It is the fact that different neurons carry the information from different frequencies that leads to our ability to discriminate frequencies. All of this discrimination is driven by the fact that different frequencies maximally stimulate different locations or places on the basilar membrane. Thus, this theory of frequency encoding is called **Place Theory [to glossary]**.

The shape of the traveling waves themselves provides convincing evidence for the place theory. But there is other important evidence that needs to be found. If the place theory is true, then it should be true that different neurons in the auditory nerve should respond to different frequencies. To examine if this is the case, several researchers have used microelectrodes and recorded from different neurons of the auditory nerve. In these studies, they stimulate the ear with different frequencies of sounds and determine the intensity of each frequency necessary to get a threshold response, which is a firing rage sufficiently above the background firing rate to be discriminated by the researcher. Figure 10.x shows the results from one such experiment (Kiang, Watanabe, Thomas, & Clark, 1962). Each of the downward peaked curves represents the recording from one nerve. Each curve has a point where it takes much less intensity to reach threshold than for other frequencies. This is termed the characteristic frequency for that neuron. These findings are very consistent with the place theory.

Play with the **Frequency of Sound** slider to see how a whole range of frequencies stimulate the basilar membrane. You can quickly summarize the height of the traveling wave by selective **Yes** in the **Show Envelope** menu. This will plot the farthest the basilar membrane moves at each point both above and below the resting point. It is called the envelope because it contains all the movement of the basilar membrane. Watch the motion of the traveling wave fit inside this outline. You can also show a guide for peak frequencies by selecting **Yes** the **Show Frequencies** menu. Get a feel for the theory and see if you can see any potential problems.

Frequency Theory. Wever and Bray (1937) picked up on a different feature of the sound wave and developed a very different theory for how frequency is encoded in the brain. Examining the shape of the sound wave with its periodic compressions and rarefactions leads to the hypothesis that neurons may fire in synchrony with the wave form. It could be entertained that neurons fire with every compression of the sound wave. Thus, in a 100 Hz tone, there are 100 compressions a second and, it follows, 100 action potentials a second in the neurons. The firing rate of each of the neurons matches the frequency of the sound stimulus giving this theory its name **Frequency Theory [to glossary]**. The frequency can be encoded directly in the firing of the neurons of the ear. Wever and Bray developed their theory decades before the tip links were discovered and a mechanism for this synchrony between sound frequency and firing rate of neurons could be discovered. Yet, the description of auditory transduction given earlier in this chapter is very consisted with this theory. The tiplinks are pulled tight and open the ion channels at only one point in the cycle of the soundwave. Thus, the signal to the auditory nerve will be in synchrony with the frequency of the sound wave.

Open **Interactive Illustration 10.x, Frequency Theory** to see how the firing rate can come to match the frequency of the sound stimulus. The screen is divided into three sections plus controls. At the top is a representation a single **Hair Cell** showing, initially, three cilia each connected by tip links. In the middle, is a graph that will indicate the motion of the basilar membrane caused by the **Sound Stimulus**. Below this graph is a graph that will indicate the firing of an auditory nerve connected to this hair cell (**Firing of Neuron**). To start the animation, click the **Start** button at the bottom of the screen. For ease of view only the motion of the cilia are indicated, by default, on the animation. If you wish to see how the motion of the cilia compare to the motion of the basilar membrane you can separately plot the titling and the up and down motion using the checkboxes next to the **Stop** button. Watching the motion of the cilia, you can observe how it follows the motion of the basilar membrane. When the cilia are pushed, in this case to the left, they spread out and the tip links are pulled tight and when completely tight, the ion channels in the cilia are opened, indicated both by text and a sound. When the ion channels are opened, the change in the voltage inside the hair cell can trigger the neuron to fire as shown on the **Firing of Neuron** graph. To really observe how the firing rate of the neuron is tied to the stimulus frequency, use the **Frequency of Sound** slider at the bottom of the screen and reduce the sound frequency to 250 Hz. At this much lower sound frequency, the rate of action potentials still matches the rate of opening of the ion changes which matches the stimulus frequency. Now raise the frequency to 1000 Hz. The firing rate is still tied to the opening of the ion channels but you might notice that some of the cycles of the sound are missed. This observation is discussed next.

There is one grave difficulty with this theory however, as described so far. Neurons have a maximum firing rate of about 1000 Hz (the animation does better by far than our actual neurons) yet our hearing goes up to about 20,000 Hz. Wever and Bray (1937) were aware of this difficulty and developed a part of frequency theory to deal with this problem. The auditory nerve is not made up of one neuron but thousands. If the encoding of frequency is not the responsibility of one neuron alone but many, it is possible to encode much higher frequencies than 1000 Hz. This part of the theory is called the **volley principle [to glossary]** because the sound is shared or volleyed across several neurons. Open **Interactive Illustration 10.x, Volley Principle** **[link to media]**. The layout of this interactive illustration is very similar to Interactive Illustration 10.x, Frequency theory. The only differences are having 10 neurons tracked across the bottom, and several buttons to use to set the frequency of the sound. Press the **2500 Hz** button to start the animation. No single neuron can encode more that 1000 Hz so no single neuron will fire to every cycle of this 2500 Hz tone. But, if you look across all of the neurons, at lest one neuron will fire every cycle. With the volley principle, frequency theory can overcome the 1000 Hz limit of the firing rate of the neurons. If we think of the firing rate collectively across many neurons, then the nervous system can keep up with faster frequency sounds.

What is needed now is some evidence that neurons collectively do actually fire collectively with the sound stimulus frequency. To measure this, a researcher puts an electrode in the auditory nerve but not penetrating any single neuron. The goal is to record from several neurons at once. The stimulus is presented and repeatedly so and the number of action potentials at each moment after the stimulus is started is recorded. The resulting graph is called the post-stimulus histogram. Looking at 10.x, it is apparent that most of the action potentials are grouped together and occur at periodic intervals (Rose, Brugge, Anderson, & Hind, 1968). These intervals happen to match the frequency of the sound providing strong evidence in favor of frequency theory.

Reconciliation. At this point there is a situation where there are two theories for frequency encoding. Both theories have evidence that support their claims. It turns out that both theories play a role in frequency encoding. Frequency theory works better for the lower range of frequencies, say less than 5000 Hz. If you open **Interactive Illustration 10.x, Place Theory and Frequency Ranges [link to media NOT DONE]**, you will see a figure similar to what was seen above but now you can use a much larger range of frequencies. As you use lower and lower frequencies, the peak of the point where the basilar membrane is maximally stimulated gets wider and when the frequency gets low enough, the entire basilar membrane vibrates together with no peak. Thus, place theory is not an effective means of encoding lower frequencies. As the frequencies get higher, frequency theory becomes less effective and place theory works best.

a. The Effect of Timbre [HERE OR ELSEWHERE?]

[NEED TO ADD COMPARATIVE STUFF]

Summary

This chapter described the nature of the sound stimulus. It then described the anatomy of the ear and traced the path of the sound stimulus from outer ear through the middle ear to the inner ear where transduction takes place. Transduction is done by the hair cells which lie on the basilar membrane of the cochlea. Transduction is accomplished by the bending of the cilia of the hair cells which pull on tiplinks and open ion channels in the hair cells. One of the interesting features of hearing is the ability to hear the component frequencies of a sound stimulus. This is known as Ohm's acoustical law. The chapter then proceeded to discuss the encoding of loudness and frequency. Loudness, like brightness, is perceived such that loudness changes more slowly than intensity. Frequency encoding is accomplished by two theories: place theory and frequency theory. Place theory finds that different frequency sound stimuli stimulate different places on the basilar membrane to the greatest extent, which leads to different neurons responding to different frequencies. Frequency theory is where the frequency of the sound stimulus is directly encoded in the frequency of the firing of neurons.

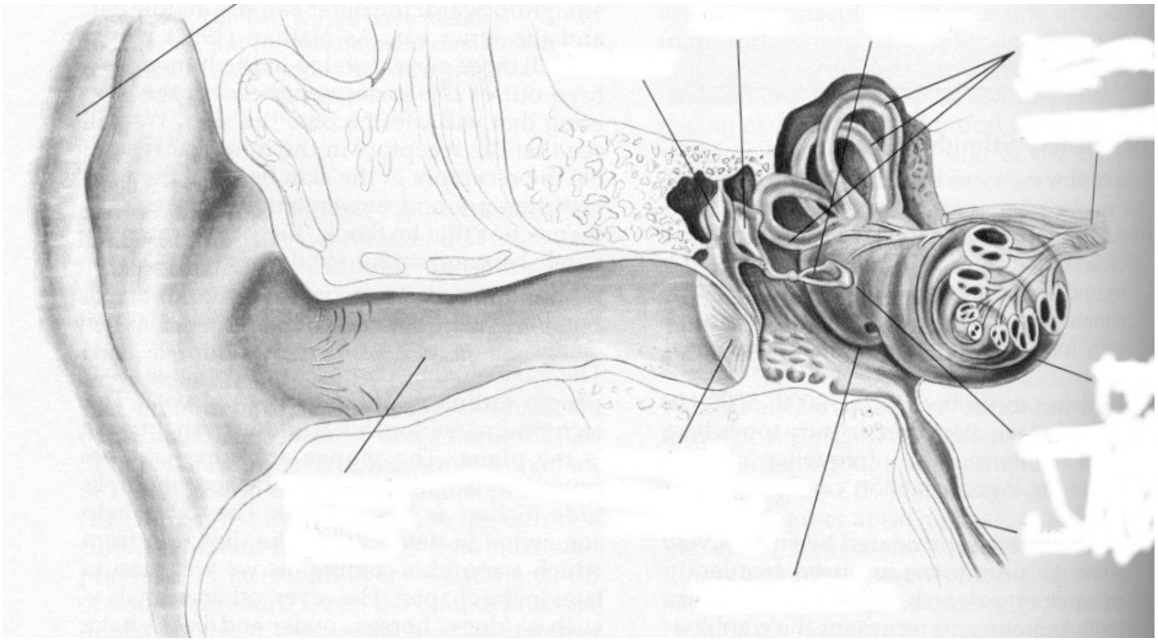

Figure 10.x. The Ear

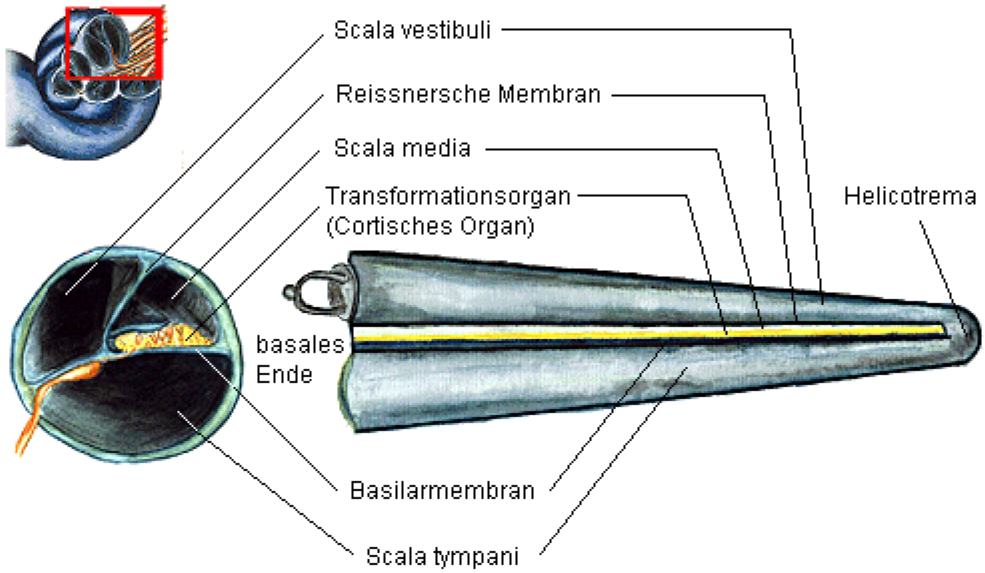

Figure 10.x. Side view of the cochlea

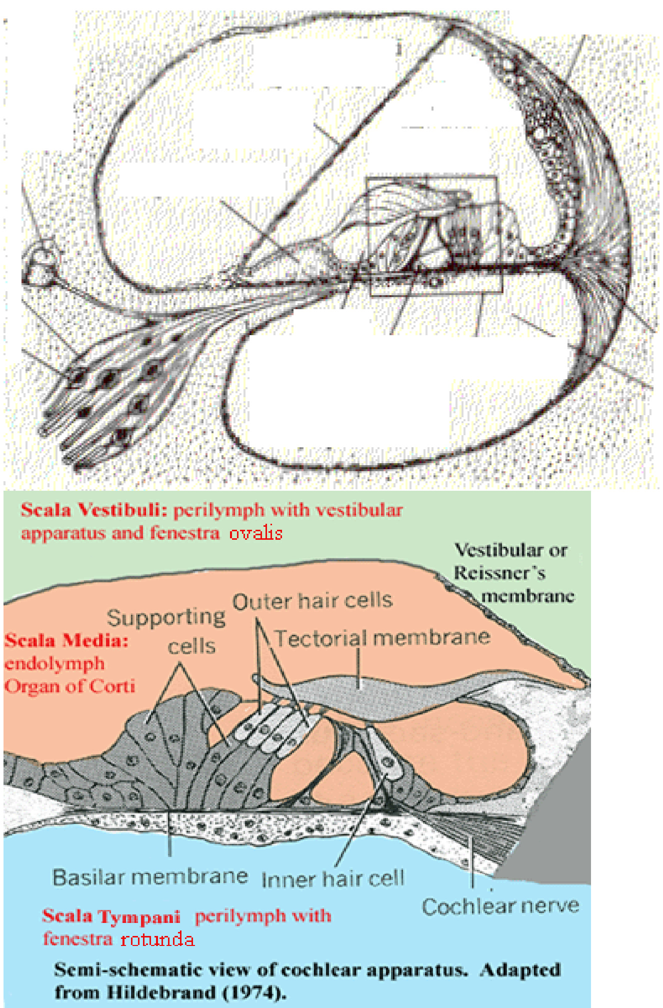

Figure 10.x Cross section of the cochlea showing the Organ of Corti

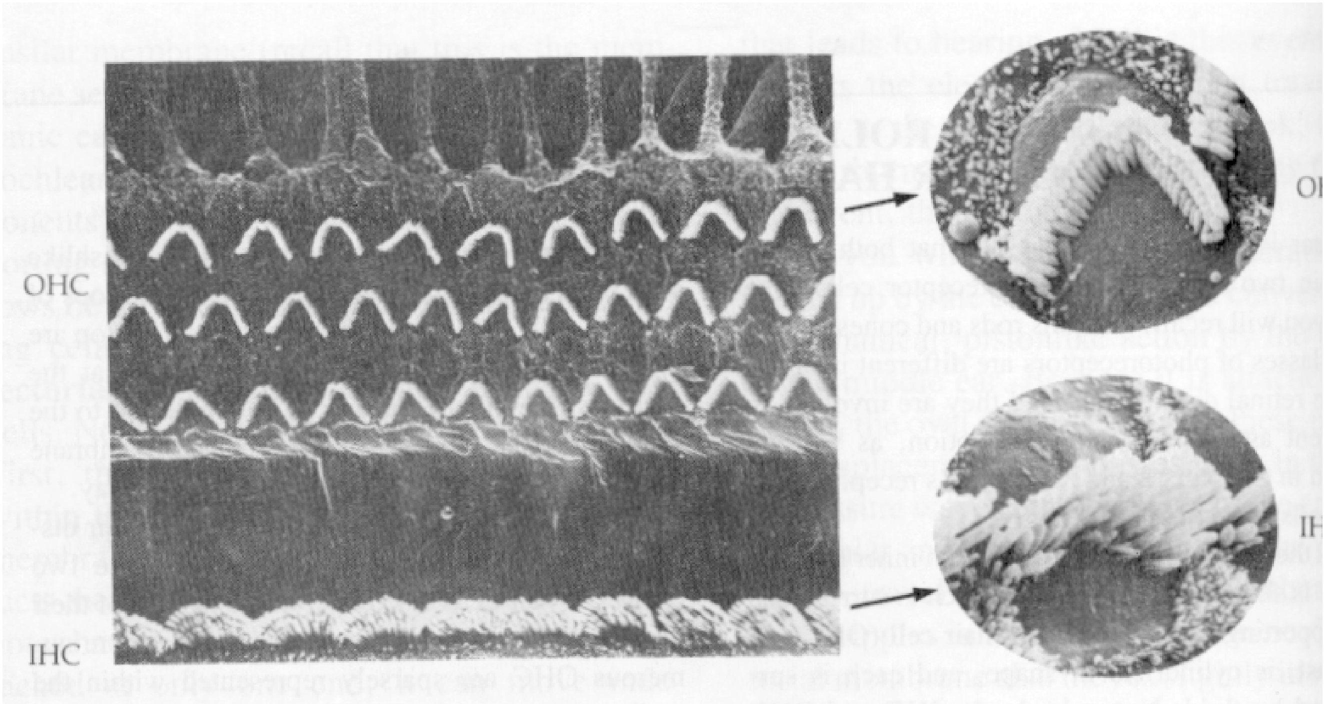

Figure 10.x. Scanning Electron Microscope image of the hair cells

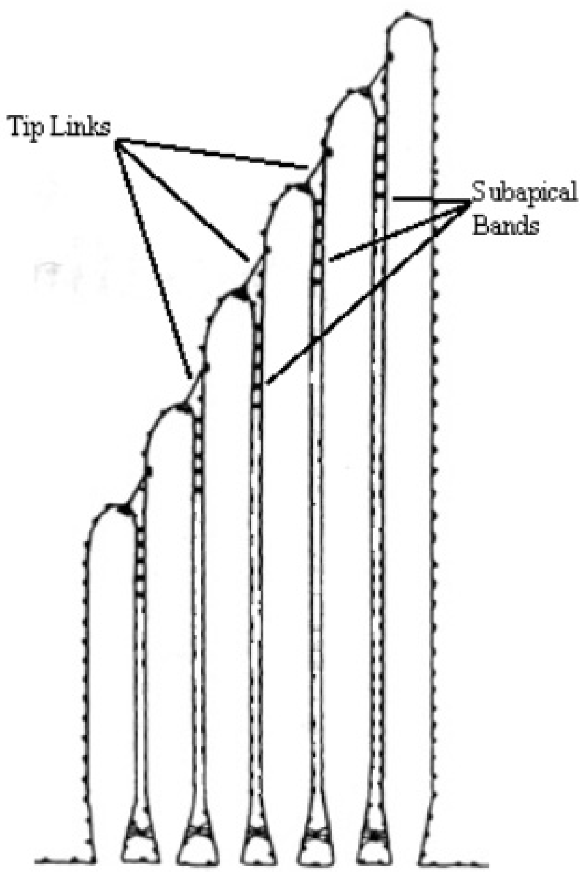

Figure 10.x, The hair cell connections for transduction.

C

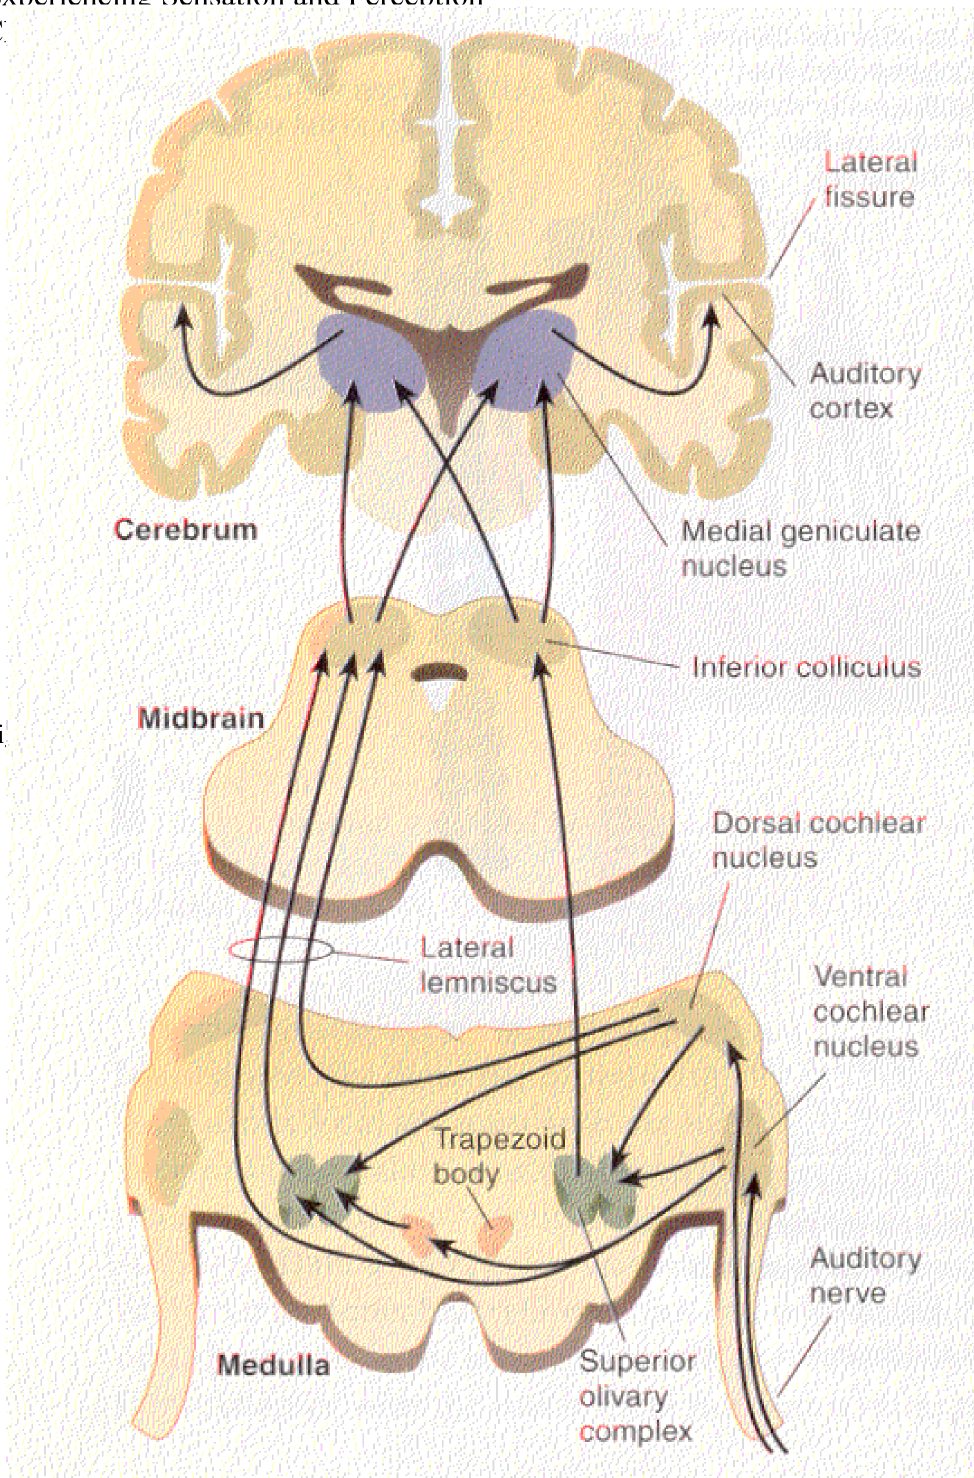
Lateral
fissure
Auditory
cortex
Cerebrum
Medial geniculate
nucleus
Inferior colliculus
Midbrain
Dorsal cochlear
nucleus
Lateral
lemniscus
Ventral
cochlear
nucleus
Trapezoid
body
Auditory
nerve
Medulla
Superior
olivary
complex

Fi

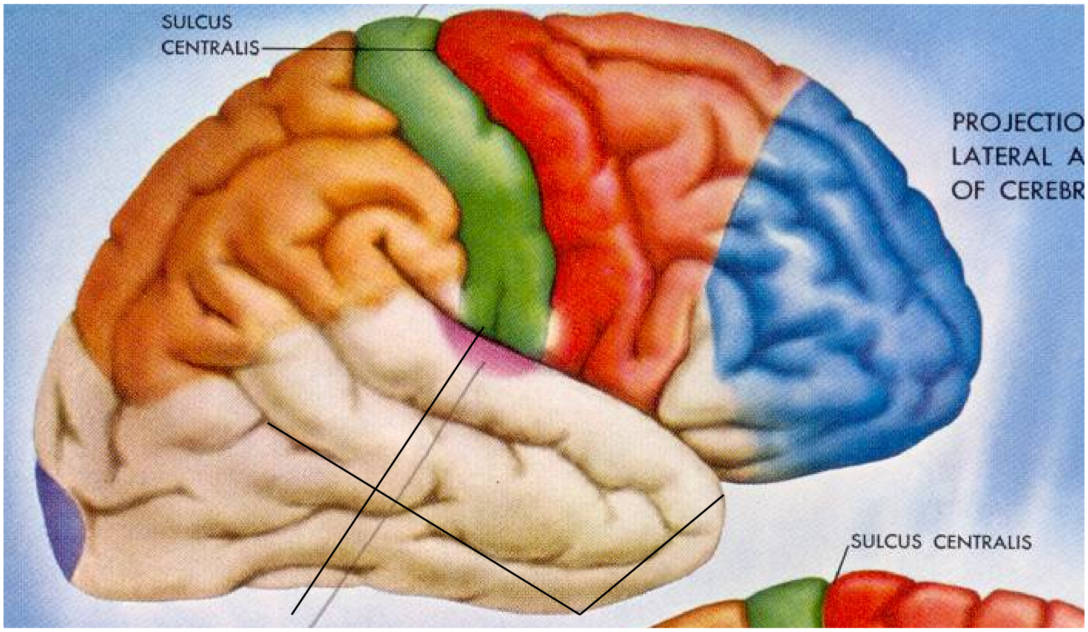

Auditory Cortex Temporal Lobe

Figure 10.x. The auditory cortex from the side of the brain.

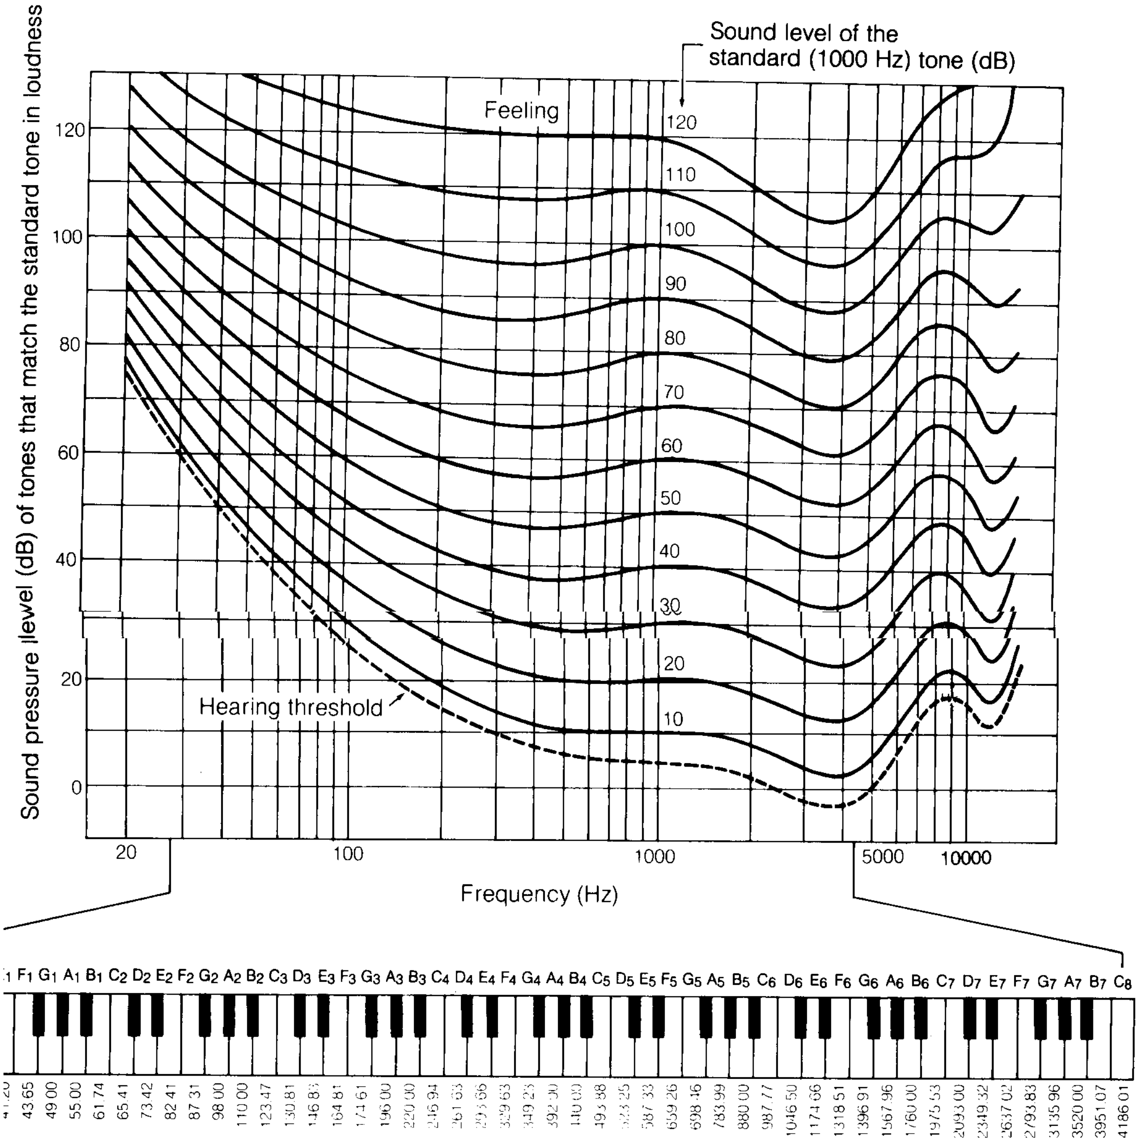

Figure 10.x. Frequency response of the ear. [This figure would only have the bottom dashed line at this point as it is just the threshold.]

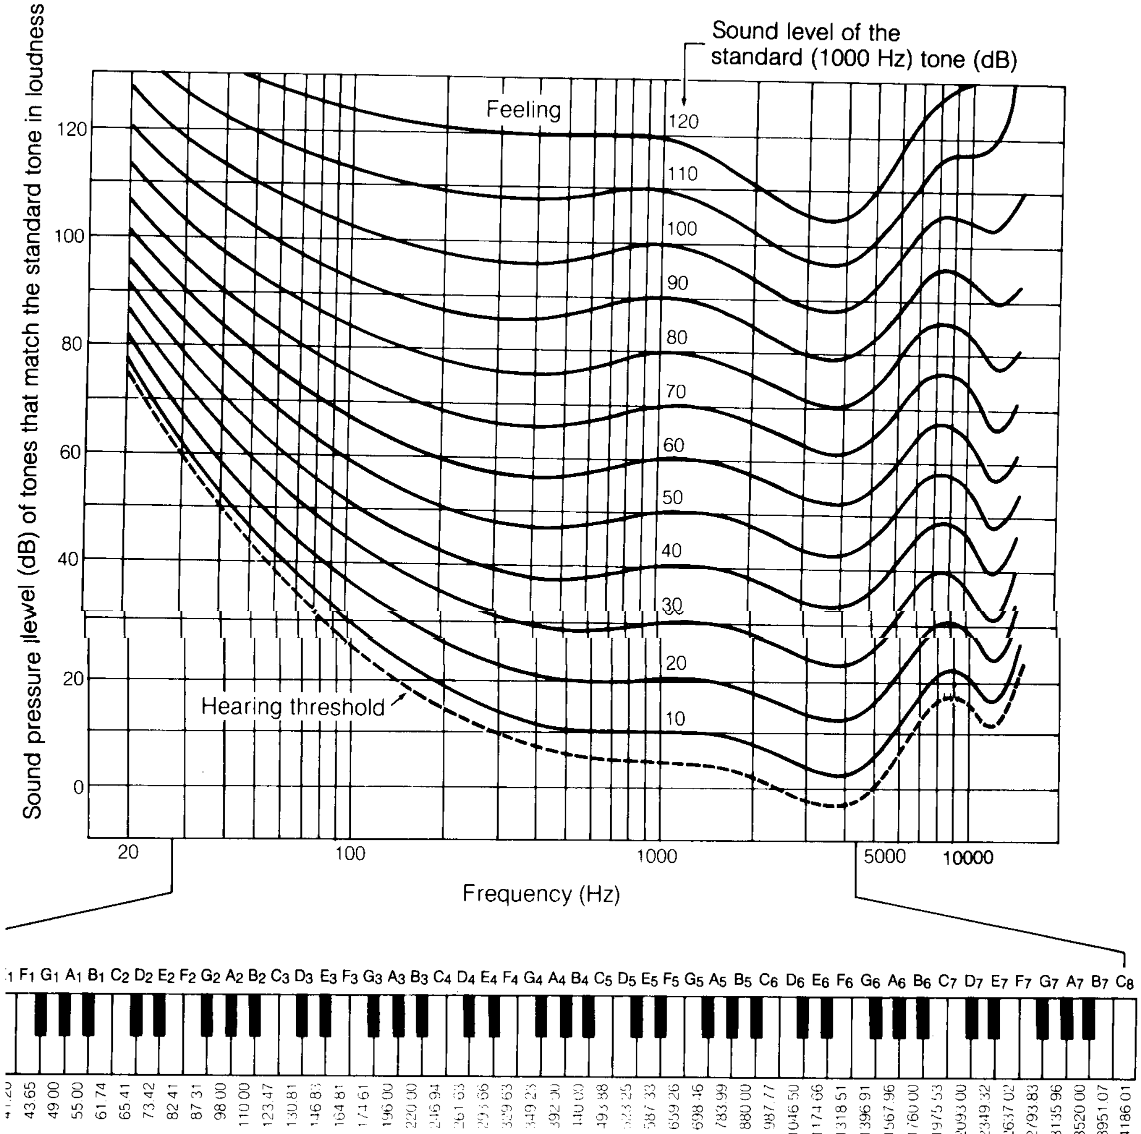

Figure 10.x. An equal loudness contour. [this graph should only have the dotted line and the line labeled 10 or this might be turned into an interactive figure.]

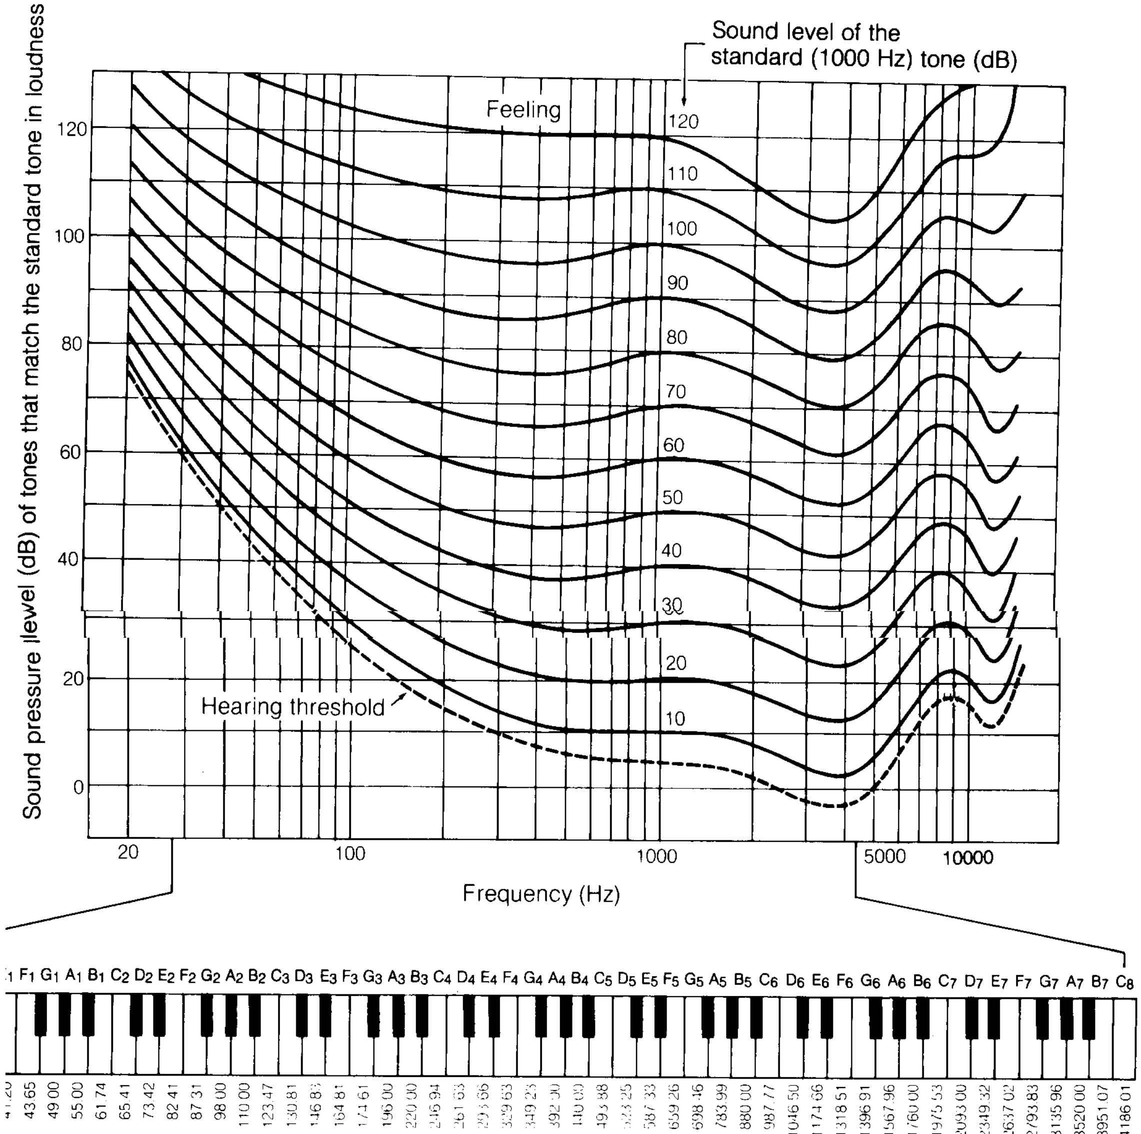

Figure 10.x. The results of the Fletcher and Munson Experiment. Each line is an equal loudness contour. [now the whole figure needs to be shown.]

Figure 10.x. Tuning curves in the auditory nerve.

Figure 10.x. Inter-spike interval histograms (Find a good figure