

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high quality educational resources for free. To make a donation or to view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at [ocw.mit.edu](https://ocw.mit.edu).

**PROFESSOR:** Last time we were talking about the hair cells. And there's a picture of a hair cell here. And what did we do about the hair cells?

We talked about the two types, inner hair cells-- plain old, ordinary receptor cells. And we talked about outer hair cells, which have this wonderful characteristic of electromotility. When their hair bundle is bent back and forth, their internal potential changes. When they're depolarized, the cell shortens. And somehow this mechanical shortening adds to the vibrations of the organ of Corti and basilar membrane that were set up by sound. And it amplifies those vibrations so that the inner hair cells then are responding to an amplified vibration.

And the outer hair cells are then dubbed by the name of the cochlear amplifier. Without the amplification, you lose 40 to 60 dB of hearing, which is a big amount. A large hearing loss would result if you didn't have outer hair cells or without their electromotility, which were shown by deleting the gene for Prestin and testing the knockout mouse, which had a 40 to 60 dB hearing loss. So questions about that?

Before we talked about the two hair cells, we talked about vibration in the cochlear, what of tuning curve is. In that case for tuning of the basilar membrane, a particular point along the cochlear-- so you could bring your measurement device to one particular place and measure its tuning in response to sounds of different frequencies, and show that a single place vibrates very nicely to a particular frequency. That is you don't have to send in very much sound into the ear. But if you go off that particular frequency, you have to boost the sound a lot to get that one place to vibrate.

And we had the example of the vibration patterns that low frequencies stimulated

the cochlear apex the most, way up near the top of the snail shell. And middle frequencies stimulated the middle. And high frequencies stimulated the basal part. We'll be talking a lot more about that frequency organization along the cochlear today, when we talk about auditory nerve fibers.

So here's a roadmap for today. We're going to concentrate on the auditory nerve. And I just put down some numbers so I wouldn't forget to tell you how many auditory nerve fibers there are.

There are approximately 30,000 auditory nerve fibers in humans. So that means in your left ear you have 30,000. And in your right ear you have 30,000 sending messages from the ear to the brain. So that's a pretty hefty number, right?

How many optic nerve fibers do you have or does a primate have? I'm sure Dr. Schiller went over that number. We're pretty visual on animals. So our sense of vision is well developed. So how many nerve fibers go from the retina into the brain compared to this number? Anybody remember?

Well, that's a good number to remember. It turns out there about 1 million optic nerve fibers from the retina into the brain. And here we have 30,000. So which is the most important sense, vision or audition? Or which sense conveys messages more efficiently, should we say?

Well, obviously, primates are very visual animals. So we have a lot more nerve fibers sending messages into the brain about vision than we do audition. So I may not have given you numbers for the hair cells. In humans we have about 3,500 inner hair cells and about 12,000 outer hair cells per cochlea. OK, so those are the numbers.

So today, we'll talk about the two types of nerve fibers. As we have two types of hair cells, we have two types of nerve fibers. We'll talk about tuning curves now for the responses of auditory nerve fibers. And we'll talk about tonotopic organization. That is organization of frequency to place within the cochlea, which is one of the codes for sound frequency. How do we know we're listening to 1,000 Hertz and not 2,000

Hertz by which place along the cochlea and which group of auditory nerve fibers is responding.

Then we'll get away from auditory nerve and have some listening demonstrations. We'll see how good we are at discriminating two frequencies that are very close together. And we'll talk about some tuning curves that are based on psychophysical measure. That is listening. You can take some tuning curves by just a human listener.

Then we'll get back to auditory nerve and talk about a different code for sound frequency. That is the temporal code for sound frequency, which involves a phenomenon called phase locking of the auditory nerve. Then we'll talk about how that's very important in your listening to musical intervals. And the most important musical interval is the octave, so we'll have a demonstration of an octave.

OK, so one of my problems as I pass by the MIT Coop on the way to class, and I always buy something. So I did a reading last week. And so we'll have a little reading from this book, *I Am Malala*. She was the girl who was shot right and recovered and was a candidate for the Nobel Peace Prize. Maybe next year she'll get the Peace Prize.

I haven't read this. I just picked it up a few minutes ago. But I went straight to the section about her surgery. So she was shot in the head on one side.

And she said, "While I was in surgery--" this is after recovery. This is a further surgery she underwent. "While I was in surgery, Mr. Irving, the surgeon who had repaired my nerve--" that's her facial nerve-- "also had a solution for my damaged left ear drum. He put a small electronic device called a cochlear implant inside my head near the ear, and told me that in a month they would fit the external part on my head. And then I should be able to hear from that ear."

OK, so the cochlear implant is a device that stimulates the auditory nerve fibers. And in a person who's had a gunshot wound-- either because of the loud sound or the mechanical trauma to the ear or temporal bone-- possibly the hair cells are

damaged or are completely missing. And the auditory nerve fibers remain. The person is deaf without the hair cells.

But the device called the cochlear implant can be inserted inside this person's cochlear to stimulate the auditory nerve. And we'll have a discussion of the cochlear implant next week when we have a demonstrator come to class who's deaf. And she'll show you about her implant.

But we do need to know a lot about the auditory nerve response before we can really think about what is the good coding strategy for cochlear implant. That is how do we take the sound information and translate it into the shocks that are provided by the cochlear implant electrodes that stimulate the nerve fibers. Because little electric currents in the cochlear implant are made to stimulate the auditory nerve fibers that can then send messages to the brain. So it's just a little motivator for what's important about auditory nerve code.

So we'll start out today with the hair cells. And these are the auditory nerves here. One thing that's interesting about vision and audition is the look of the synapse between the hair cell and the nerve fiber, and between the photoreceptor-- you have rods and cones in the retina. And they have associated nerve terminals here.

And these are electron micrographs, taken with a very high powered electron microscope that looks at the synapse between the photoreceptor up here or the hair cell up here and the associated nerve terminal down here, or the associated either horizontal cell or bipolar cell down here. So in each case, you have obviously the synapse here is a little gap. And you have synaptic vesicles that contain the neurotransmitter. And they're indicated here in the photoreceptor. SV is the vesicle. And inside that vesicle is the neurotransmitter.

When the receptor cell depolarizes, these synaptic vesicles fuse and release their neurotransmitter into the cleft and fire or activate their post-synaptic element. In the case of the hair cell, the auditory nerve fiber.

This structure here is called the synaptic ribbon. And it's supposed to coordinate the

release of the vesicles. And they call it a ribbon in the hair cell here, even though it looks like a big which ball. It doesn't look like a ribbon at all.

But it's called a ribbon, because it has the same molecular basis. It has a lot of interesting proteins and mechanisms to coordinate the release of these neurotransmitter vesicles, which presumably are synthesized up here in the cytoplasm and are brought down to the ribbon and coordinated and released at the hair cell to nerve fiber synapse. So I just wanted to show you the look of the synapse in the electron microscope. So that's what it looks like.

And the next slide here is this schematic of the two types of hair cells, inner hair cells and the three rows of outer hair cells, and their associated nerve fibers. And I think I mentioned last time that almost all of the nerve fibers, the ones that are sending messages to the brain at least, are associated with the inner hair cells. So you can see how many individual terminals there are-- as many as 20 on a single inner hair cell.

By contrast, the outer hair cells-- you can see, well, this one has three of them. But they're all coming from the same fiber, which also innervates the neighboring hair cells. So there are very few of these so-called type two auditory nerve fibers.

Here are the numbers. So this total is in cats. Cats have more nerve fibers than humans, a total of maybe 50,000. About 45,000 of them are the type ones, associated with inner hair cells, and only 5,000 or the type twos associated with outer hair cells. So you can see by this ratio then that most of the information is being sent into the brain by the type one fibers, sending messages from the inner hair cells.

Those axons of the type one fibers are thick. They have a myelin covering, compared to the type two fibers, which are very thin and they're unmyelinated. And actually, one of the very interesting unknown facts about the auditory system is that as far as we know, no recordings have ever been made to sample the type two responses to sound.

Do they respond to different frequencies? Are they widely tuned? Narrowly tuned? We don't know that at all. And it turns out that it's just very difficult to sample from such thin axons as you find in the type two fibers.

So I actually have a grant submitted to the National Institute of Health to use a special type of electrodes to record from the type twos. I think it's being reviewed next week. And I hope it gets funded because then maybe I'll figure out this mystery. But it will be challenging not only because they're thin, but because there are fewer of them.

So when I talk about auditory nerve fiber recordings for this class, I'm going to be talking about the type ones. That's the only kind we know of. And here is an example tuning curve or receptive field for a type one auditory nerve fiber. Now, I think Peter Schiller probably talked about single unit recordings with micro electrodes.

So you have your nerve. It could be the optic nerve. It could be the auditory nerve, which is what we're talking about. You have a microelectrode, which is put into the nerve. And the tip of the microelectrode is very, very tiny. It could be less than 1 micrometer in diameter.

And usually the electrode is filled with a conducting solution like potassium chloride. And the pipette that's filled with a KCL comes out to a big open end. And you can stick a wire in here and run it to your amplifier and record the so-called spikes.

You guys talked about spikes, right? So you're recording the spikes, AKA action potentials, AKA impulses. And if you want to do this in a dramatic way, you send this signal also to a loud speaker and you listen to them. And maybe we'll have a demonstration at the end of the year on these. It's pretty nice to listen to that.

So you put your electrode in there. And you move it around until you have what's called a single unit. And why is it called a single unit? Well, in the old days, people didn't know what was being recorded.

Is it a cell body? Is it a nerve axon? Is it the dendrite? What is it ?

All they knew is that coming out of the amplifier, they saw this spike. And that's what's plotted here. These are a bunch of spikes.

And it's called a single unit, because most of the time when you get one of these recordings, the spikes look all the same. But every now and then you get a recording that looks like this. And this is interpreted as being fiber or axon number one. Here's another number one.

And this is a second fiber that's nearby. But it's a different one. Maybe there were actually two fibers right next to each other. And you could record both of them. That's very unusual.

More commonly, you just have a recording from one single unit. And the interpretation is you are sampling from just one auditory nerve fiber out of a total of 40,000. Is that clear?

So such experiments are done in the auditory nerve. In this case, I think the experimental animal was a Guinea pig. And in this case, it's recordings from a chinchilla auditory nerve.

So what's the stimulus? Well, this is a plot of sound frequency, sound frequency in kilohertz. And this axis, on the y-axis, is sound pressure level. So this is how loud it is, if you will.

And at a very low or soft tone level, if this frequency is swept from low to high frequencies, there was hardly any spikes coming from that single unit. But if you boosted the level up a little bit. And you came to a frequency that it was about 10 kilohertz, there were a bunch of spikes produced by that single unit.

Then if you boosted the level up so it was a moderate level, there were spikes anywhere from 8 kilohertz up to 11 kilohertz. All that band of frequencies caused a response. Then at the highest level, everything caused a response, from the lowest frequencies up to about 12 kilohertz, and nothing above.

What's this activity out here? I said nothing above and nothing over here. Well, there's some spontaneous firing. So even if you turn the sound completely off, these nerve fibers have a little bit of activity. They fire some impulses. There's an ongoing thing.

If you outlined this response area with a line-- that line is the border, say, between spontaneous firing or no firing and a response. So inside of the receptive area there's a response. And outside there's nothing. Those lines are called tuning curves.

And here are a bunch of tuning curves from a chinchilla. And there are one, two, three, four, five, six different tuning curves. So what the experiment did was they moved the electrode in and got one single unit.

And then they moved the electrode, let's say deeper into the nerve. And now they sampled a different neuron, a different single unit. OK, maybe got this tuning curve. Then they went deeper and sampled from this one and this one and this one and this one.

And the idea that it's a different one-- well, the response is different. But also, as you move the electrode, you lost the single unit, number one. And you've maybe put it deeper, a millimeter or so. It's a huge distance.

And you've got a new unit. The action potentials probably look different. That's a second.

OK, so these are tuning curves there then from six different single units. And each of them comes down to a pretty nice tip. And if you take that tip and the very lowest sound level they're caused a response and extrapolate that to the x-axis, you get a frequency. And that frequency is called the CF, or characteristic frequency.

OK, so CF is a very important term. You should know that the CF is the very tip of the tuning curve. And the CF is different from frequency.

Frequency is whatever you want to dial in with your sound oscillator. But CF is a

particular characteristic of a neuron, in this case an auditory nerve fiber, that you're recording from. And it's a characteristic that it has that you measured from it.

Many of these tuning curves, in addition to having a CF and a so-called tip region, also have a tail. And in this very high CF neuron, the tail goes like this. And then there's actually, I think, something that I dashed in here, a dashed line here. And the tail continues way down here.

These experiments didn't want to boost the sound level to get all the tail above 80 dBs because of possible damage. If you crank up too much sound-- just like you get a gunshot to the head is a very loud sound-- it can cause damage to the hair cells. They didn't want to do that. But you could see the tail of this response area. It's a nice tip and a nice tail.

OK, now, right away we have a beautiful potential code for sound frequency. How do I know I'm listening to 8 kilohertz? Well, this nerve fiber responds very nicely, lots of action potentials.

How do I know I'm listening to 1 kilohertz? Well, that same nerve fiber might respond. But I have to get the sound level to very loud level, like 80 dBs ATL. But these other guys over here with CFs of 1 kilohertz would respond at a very low sound.

So then we have a code of which fiber you're listening to tells you which frequency you're listening to. It's very important. You judge an instrument, like a violin, by its combination of frequencies. A guitar has a different recombination frequencies. Male speakers generally have deeper voices than female speakers, deeper meaning more low frequencies. Female and children's voices are higher in frequency. So frequency is essential for you to identify what sound stimulus you are listening.

Why do we call this a place code for sound frequency? Well, as we talked about before, different parts of the cochlea respond to different frequencies. Here is a beautiful example of the place map for auditory nerve fibers.

And in this case, microelectrode recordings are done as we described before. But instead of just a plain old potassium chloride solution in the microelectrode, it's filled with a substance called a neural tracer. What are examples of neural tracers? Has anybody played around with neural tracers before? Give me some examples of chemicals that are neural tracers. Anybody?

This one is a funny name horseradish peroxidase, abbreviated HRP. Another one is biocytin, OK, biotinylated dextran amine, PDA. There's millions of them, Lucifer yellow. You can tell that I'm a tracer kind of guy. I use tracers all the time in my experiments.

So what you do with these neural tracers, it's convenient if they are charged. For example, horseradish peroxidase, this is a positive charge. And you can apply positive current to the pipette up here. That's going to tend to force positive charge out the tip of the electrode. You can expel a positive ion out the tip by this technique, which is called iontophoresis. And if it happens that your tip is close to or ideally inside an axon, some of that HRP is going to come out the tip of the electrode and go into the axon.

And why did we pick HRP? Because it's picked up by chemical transport systems that transport things along axons. And there are several of these. There's fast axonal transport. There's slow. There's medium.

There's a whole bunch of systems, because this axon is coming from a cell here. It's connected to the cell body. In the cell body you make things like neurotransmitter, because that's where you can make protein. And that neurotransmitter has to get down to the tip of the axon, which in the case of the auditory nerve is in the cochlear nucleus of the brain. So there are all these transport systems transporting things.

And it just turns out that some chemicals are picked up by them. HRP is one of them. When you iontophorese HRP into a nerve fiber, it's transported to all parts of the nerve fiber, including to the cell and including out to the tip of the nerve fiber on the hair cell.

So here is an example of iontophoretically labeled nerve fibers. And there's five or six of them. The recording site was here in the auditory nerve.

This is a diagram of the cochlea. This is the so-called Schwann glial border, which defines the periphery and the brain. So this would be the brain. So these are the nerve fibers going into the brain.

They were recorded in the auditory nerve. And you can trace them out into the periphery. Right here is the cell body of the auditory nerve fiber.

Every neuron has a cell body. Most neurons have axons. The axon was what was recorded.

And the auditory nerve neuron has a cell body. And it also has a peripheral axon that goes out to the periphery and contacts an inner hair cell. As we saw before, these are type one auditory nerve fibers going to inner hair cells. And it contacts usually one inner hair cell.

Now, you can know exactly where that auditory nerve fiber started out by tracing it and by tracing the base of the cochlea through the spiral and all the way up to the apex. So starting at the base to the apex-- so that's 100% distance, let's say. And if this were halfway between the base and the apex, that would be the 50% distance place.

This guy ending up near the apex might be 80% distance from the base to the apex. OK, does everybody see how I can make that mapping? These sausages here are the outlines of the ganglion, the spiral ganglion, where the cell bodies are of the auditory nerve fibers.

So what good is that mapping? Well, before we put the HRP in, we measured the tuning curve. And we got the CF from the tuning curve.

So we measured the CF. We injected the neurotransmitter to label the auditory nerve fiber. And we reconstructed where the labeled ending of the auditory nerve fiber contacted its inner hair cell.

Why did we do this for five of these? Well, in the ultimate experiment, you just do it for one. But if you're getting good at reconstructing the mapping, you can tell it should be about the 50% place. And you go and find it's 51%. You know that fiber was different than the one up there.

Then you make your mapping-- characteristic frequency to position of innervation along the cochlea. And here is the mapping. These are the CFs.

And this is the percent distance along the cochlea from the base. So 0% distance from the base would be the extreme base. 100% distance would be the extreme apex.

And you can see this beautiful mapping of CF to position, almost a straight line, until you get to the lowest CFs. And this, as usual, in the auditory system, this frequency axis, this is the CF axis now. It's on a log scale. So log frequency maps to linear distance along the cochlea. Now, if the brain hears that the 50% distance auditory nerve fiber is responding and no other auditory nerve fiber is responding, it knows it's listening to a 3 kilohertz frequency.

Place to frequency mapping is tonotopic. I said that opposite. Frequency to place is tonotopic. So this is a tonotopic mapping-- frequency to place, tonotopic.

And why is that important? Well, it happens in the cochlea. It happens in the auditory nerve. It happens in the cochlear nucleus of the brain. It happens in almost all the auditory centers in the entire brain, all the way up to the cortex.

You have neurons or fibers responding to low CFs over here in the brain. And if you move your electrode over here, you find they're responding to mid frequencies. And if you move them over here, they're responding to high CFs.

So this organization is fundamental. It starts at the receptor level in the cochlea. It's conveyed by the nerve into the cochlear nucleus. And you have these beautiful frequency-- they're actually CF organizations in the brain. So the place code for sound frequency presumes that each frequency stimulates a certain place along

cochlea.

And I guess, if you generalize this from the auditory system to the visual system-- if you have a particular light source, like that light over there, and my eyes are looking this way, that light is going to stimulate a particular place in my left retina and in my right retina. So you have a coding for where that light is along the place in the retina. In the auditory system, you don't have that kind of a place code. You have a place code for sound frequency. It's very different. The cochlea maps frequency.

How can we use this code? We're actually very good at distinguishing closely spaced frequencies. And here is now some psychophysical data from human listeners. We're going to get away from the auditory nerve for awhile and talk about listening studies.

Here is to graph of frequency. And on the y-axis is delta f. What's delta f? Delta f is the just noticeable difference for frequency. Of course, we're talk about sound frequency.

And how is the experiment conducted? Well, you have your listener. Your listener is listening to sound.

And you give them a 1 kilohertz sound. And then you give them a 2 kilohertz sound. The experimenter says, does it sound the same or different? Ah, completely different.

OK, 1 kilohertz sound and a 1,100 kilohertz sound, same or different? Ah completely different. OK, 1,000 hertz sound and 1,010 hertz sound. Ah, it's different.

1,000 hertz and a 1,002 hertz sound? I'm not so sure. Give it to me again.

OK, 1,000 hertz sound, 1,002 hertz? Eh, it's just a little bit different. 1,000 hertz sound and a 1,001 hertz sound? Same. OK, so that's the experiment.

So we have the graph here for the just noticeable difference in frequency, as a function of frequency. And at 1,000 hertz-- that's right in the middle of your hearing range-- the delta f-- well, it's hard to read that axis-- the delta f is about 1 or 2 hertz.

So 1,000 vs 1,002 hertz is just barely distinguishable for human listeners.

You can do that experimental a little bit differently. Instead of giving two tones, you can give one tone and vary its frequency a little. And that's kind of a pleasing sound. Does everybody know what a vibrato is on a stringed instrument?

That's a plain A. But if you vibrate it a little-- that's the frequencies going back and forth. Everybody could hear that vibrato right? Even though I'm changing the frequency just a tiny bit.

You could do the experiment by vibrating the frequency just one single frequency. Is it vibrating? Or is it not vibrating? And you get about the same result. That's what the second graph is.

People who are tone deaf, not proficient music, don't have any hearing problems are almost always able to distinguish frequencies with a little bit of training. The training is now here's the task, that type of training.

OK, so I have a demonstration here. And we can listen to this and see how good you guys are-- you know, naive, untrained listeners-- and see if we're good at distinguishing frequency. So the demonstration is a little bit complicated. So I'll go through it.

It's going to give you 1,000 hertz, standard, middle of your range hearing frequency. And it's going to give you a bunch of different groups. I'm going to go through these slowly.

And in each group, we have 1,000 one hertz, and 1,000 hertz plus  $\Delta f$ . OK,  $\Delta f$  for group one is 10 hertz, big frequency space. And what you're going to listen to is A, B, A, A, where is  $f$ -- 1,000 hertz-- and  $f$  plus  $\Delta f$ -- 1,010 hertz.

And B will be first, 1,010 hertz and then 1,000 hertz. Then A, 1,000 hertz, 1,010 hertz; and another A, 1,000 hertz, 1,010 hertz, just to give you a bunch of different examples.

Then group two, delta f will be a little bit harder, 9 hertz, OK, so on and so forth, down to group 10, which will be delta f of 1 hertz. Or seeing if we can distinguish 1,000 and 1,001 hertz. OK, so let's listen to this.

**MAN IN AUDIO:** Frequency difference file for J and D. You will hear 10 groups of four tone pairs. In each group, there is a small frequency difference between the tones of the pairs, which decreases in each successive group.

[BEEPING OF TONE PAIRS]

**PROFESSOR:** That's group one. Here's group two.

[BEEPING OF TONE PAIRS]

**PROFESSOR:** OK, could everybody do the big interval, delta f equals 10? Raise your hand if you could do that. Most-- some people can. I-- it's not problem. OK, how about your limits for people who could do it. When did you--

**AUDIENCE:** I heard eight.

**PROFESSOR:** About eight, OK. And what was your limit going down?

**AUDIENCE:** Nine.

**PROFESSOR:** Group nine or delta f? OK, so delta f of 2. I cut out about between two and three. Well, for those of us who could do it, without any training at all, you get to what the best results are-- people who have done this for days and days and practice. And this is not an ideal listening room. There's a lot of fan noise. There's some distractions too.

Ideally, you'd be in a completely quiet environment, perhaps wearing headphones. But it works pretty well. I'm not sure what it says about people can't do it. And there certainly are. So I don't know if you should have your hearing tested or whatever. But for those of us who could do it, you get quickly to the best possible results.

So you could do a calculation then on these. We know what  $\Delta f$  is. Let's say it's 2 hertz at 1,000 hertz. And let's go back to our mapping experiment.

So here's 1,000 hertz CF. Let's say we're listening at the CF. And we're moving from 1,000 to 1,002 hertz, the best possible psychophysical performance.

We can go up along this cochlear frequency map and say, well, what percent distance did we move from the 1,000 hertz point to the 1,002 hertz point? And I don't know where it is. Well, it's about the 70% distance place in this animal. This is a cat, of course. You can't do these kinds of studies in humans. You could map it out in human.

It turns out that if you know how many inner hair cells there are-- we had that number before along the base to apex spiral-- and you know the distance you're moving, it turns out you can make the calculation, the best possible performance. You're moving from one inner hair cell to its neighbor. So it's a very, very small increment along the cochlear spiral you're moving.

That increment is associated with the best possible psychophysical performance in terms of frequency distinction. OK, that's the cochlear frequency map. OK, any questions about that?

Now let's go back to the auditory nerve and talk more about coding for sound frequency. So far, we've just been exploring single tone response areas. So now let's make the stimulus a little bit more advanced and talk about coding for two tones.

What happens when you have two tones? Well, here is a tuning curve, plotted with open symbols here, for the kind of tuning curve with one tone we had before. So everything within this white area is excitatory. You put a frequency of 7 kilohertz in at 40 dB SPL. And the neuron is going to fire all kinds of action potentials.

Now, let's put in a tone right at this triangle, called the probe tone. It's usually right at the CF. And it's above the threshold. In this case, it looks like it's about 25 dB. And it gets the neuron responding. You put that probe tone in the neuron is going to

fire some action potentials.

And keep that probe tone in so the neuron is firing action potentials. And put a second tone in. And the second tone is often outside the response areas.

And it turns out that anywhere in this shaded area above the CF or below the CF, a second tone, as is illustrated here, will decrease the response to the probe tone in a dramatic fashion. Then, when you turn off this second tone-- sometimes called a suppressing tone-- the original activity comes back. And this phenomenon is called two-tone suppression.

At first, it was called two-tone inhibition. People thought, oh, OK, there's another nearby neighbor nerve fiber that's inhibiting this first one. And they looked in the cochlea and there weren't any inhibitory synapses.

OK, so that was a problem. They started calling it suppression. And they actually ended up finding it in the movement of the basilar membrane. So it's just something about the vibration pattern of the cochlea that causes the movement of the membranes to be diminished by a second tone on either side of the first.

Now, why do I bring this up? Well, it's kind of interesting in a number of contexts. Two-tone suppression might be a form of gain control. If you just had the excitatory tuning curve, and you started listening in a restaurant where everybody was talking and there was a lot of sound, all your auditory nerve fibers might be discharging at their maximal rates.

And you wouldn't be able to tell the interesting conversation your two neighbors were having. You wouldn't be able to eavesdrop. You wouldn't be having a conversation yourself.

So two-tone suppression is a form of gain control, where the side bands reduce the response to the main band, so that not everything's being driven into saturation. That's one reason. And a second reason is you can actually use this in a sort of a tricky psychophysical paradigm to measure the tuning of human listeners.

We obviously can't go into a human's auditory nerve with a microelectrode, although it's been done a couple times in surgery. But it's rare. It's easy to do a so-called two-tone suppression paradigm, where you have the person listen to the probe tone.

You say to the listener, here's a tone. I want you to listen to that. I'm going to put in a second tone. I ignore that. Don't worry about it. Just listen to that first probe tone. Tell me if you can hear that original probe tone.

Ah, yeah, sure I can hear. I'm going to put a second tone. Oh, I can't hear the probe tone anymore.

OK the second or side tone has suppressed the response to the probe tone. And you can use that as a measure of tuning, because these suppression areas flank the excitatory area. And so here are some results from humans in a so-called psychophysical tuning curve paradigm.

And these are a half a dozen or so tuning curves. Each one has associated with it a probe tone or a test tone. The task is listen to that test tone and tell me if you still hear it or if it's gone away.

The experimenter introduce a second tone, a so-called masker, at those frequencies and levels. And at where the line is drawn, the person who's listening to the probe tone says, I can't hear that probe tone anymore. Something happened to it. Well, two-tone suppression happened to it. The masker masked the response to the probe tone or test tone.

And look at the shapes of those tuning curves. They look like good old auditory nerve fibers. They have a CF.

The CF is right at the probe. They have a tip region. They have a tail region.

If you measure the sharpness, how wide they are, they're really sharp at high CFs. And they get a little bit broader as a CFs goes down. At high CFs they have a tip and a tail. At low CFs they look more like v-shape.

We can go back to the auditory nerve tuning curve with those in mind. And look how

similar they are. Here's high CF, tip and the tail. Low CF, just sort of a plain v ad they're wider. Human psychophysical tuning curves have that same general look.

Now, remember, this is a very different paradigm. Here there are two tones. The probe tone is one of them. And the masker or the second suppressor tone is the second one. Whereas in good old fashioned auditory nerve fiber tuning curve there was just one, the excitatory tone.

OK, so psychophysical tuning curves are obtained from humans in the following paradigm. We went over that. These tuning curves and the neural tuning curves from animals are roughly similar.

Now, what would you expect to happen to these tuning curves and the neural tuning curves if you had an outer hair cell problem? And this is kind of the classic-- oh, yeah, you can sort of pass that around-- a classic exam question. Draw a tuning curve. So you label this with frequency. This is the sound pressure level for a response-- I don't know. We can say whatever response you want to-- 10 spikes per second.

Label the CF. Here it is. This is a normal.

What's the axis here? Well, the CF might be-- the threshold might be at 0 dB. The tail comes in-- let's go to our animal tuning curve just so we get this right. Oops, pressed the wrong button.

OK, so the tip on this one it's about 20. The tail is coming in about 60. So we are starting down-- well, let's say it's 20. This is going to be 60 dBs SPL-- normal.

Draw the tuning curve in an animal where the outer hair cells are damaged. Well, you could say, there's no response. That wouldn't be quite right.

OK, remember we're-- this is the nerve fiber we're recording from, a type one. This is the inner hair cell. These are the outer hair cells.

And we're saying damage them, lesion them. You could have it in a knockout animal

where they had lost their Preston. OK, so the cochlear amplifier is lost.

What sort of a hearing loss do you have when you lose the cochlear amplifier? 40 to 60 dB, right? Well, what's this interval? 40 dB, right?

And it turns out, when you record from a preparation in which the outer hair cells are lesioned, this is the kind of tuning curve you find when the outer hair cells are killed or lesioned-- a tip-less tuning curve. At least from these high frequencies that have a tip. And the lows they look more bowl shaped. But there's a 40 to 60 dB hearing loss. You're not deaf, but you have a greatly altered function.

How good would this function be for telling the difference between 1,000 hertz and 1,002 hertz? Not so good, right? You need a very sharply tuned function to tell or discriminate between two closely spaced frequencies.

If you have an outer hair cell problem, not only are you going to be much less sensitive, but you're not going to be so good at distinguishing between frequencies. Another way to think about it is that if there were a whole bunch of frequencies down here and your hearing aid boosted them, you wouldn't be able to listen to your characteristic frequency anymore, because these side frequencies were getting into your response area.

So these are non-selective response areas, where the normal or sharply tuned are very selective. And what are they selective for? For sound frequency.

OK, so the outer hair cells give you this big boost in sensitivity and sharp tuning of the tip. That's the cochlear amplifier part of the function.

OK, now, how could we do this? Well, recently, within the last 10 years, you can have a [? knocked ?] out animal. But in the old days, you could lesion outer hair cells by many means.

You could lesion them by loud sounds. Well, loud sounds actually end up affecting inner hair cells a little bit as well. So the preferred method of lesioning outer hair cells was with drugs. For example, kanamycin is a very good antibiotic. It kills

bacteria. Unfortunately, it's audatoxic. It kills hair cells.

And if you give it to animals in just the right dose, you can kill the outer hair cells, which for some reason-- it's not known-- are more sensitive to them. If you give them a higher dose, it will also kill the inner hair cell. But you can create animal preparations in which the outer hair cells are gone and the inner hassles are remaining, at least over a particular part of the cochlea. And from that part, you can record these tip-less tuning curves. OK, so that is mostly what I want to say about place coding for sound frequency.

And now, I want to get into the second code for sound frequency that we have, which is a temporal code that's based on the finding of temporal synchrony in the auditory nerve. This is the so-called phase-locking. Again, we're doing the same kind of experimental preparation. We stick a recording electrode in the auditory nerve. And we record from one single auditory nerve fiber.

And we measure its spikes. Each one of these little blips is a spike. The very top trace is the sound wave form. The next trace is the response of the auditory nerve fiber. And these are superimposed multiple traces. And that trace is with no stimulus. So this auditory nerve fiber is obviously very happy firing long, spontaneous activity.

Then let's turn the sound on. The top trace is on now. This is with the stimulus.

And look at how these auditory nerve fiber impulses tend to line up at a particular phase of the sound stimulus. What's phase? Well, it's just the degrees, the sine wave, as a function of time, the sound pressure-- this is sound pressure-- and it's going through 360 degrees of phase here-- 180 degrees here.

And it looks like many of the spikes are lining up around 80 degree point. So a lot of the firing is right here. Not so much firing here. Not so much firing here. And then another waveform comes along and you get some more firing about the same time.

Now, one very common misconception about phase-locking is that every time the sound wave form goes through-- in this case 80 degrees-- the fiber fires an impulse.

That's not true at all. Here is a single trace, showing excellent phase-locking.

And there's a response to the first wave form. But then the fiber takes a break and doesn't respond during the second. And it looks like it responds on the third and the fourth. But then it takes a longer break and doesn't respond at the fifth or sixth, but it responds at the seventh, and not at the eighth or ninth, then on the 10th and 11th.

So it doesn't matter. You don't have to respond in every single waveform. You can respond in one wave form and take a break for 100 waveforms, as long as when you respond, the next time it's on the same point or in the same phase in the sound wave. So typically, to get these data, you average over many hundreds or even thousands of stimulus cycles, where one complete cycle is 0 to 360 degrees.

These are plots of auditory nerve firing. So this is a firing rate across percent of total impulses. This is now a time axis. So we're just saying when does it fire along the time.

And the stimulus, I believe, here is 1,000 hertz, so it's the middle of the hearing range. And this is excellent phase-locking. If you were to quantify this-- there are many ways to quantify this-- but you could fit, for example, a Fourier series, to that. And you could plot just the fundamental of the Fourier series. And that's what's known as the synchronization coefficient. And plot it as a function of frequency.

You could make your measurements at 1,000 hertz, which is this point on the graph. You could make them at 5,000 hertz. You could make them at 500 hertz.

This synchronization coefficient ends up being between 0.8 and 0.9 for low frequencies. And then it rolls off essentially to be random firing at around 3,000 or 4,000, certainly by 5,000 hertz. So this behavior, this phase-locking goes away toward the high end of our hearing range. It just means that the auditory nerve can no longer synchronize at very high frequency.

So what's going on here? The auditory nerve fiber is getting its messages from the hair cell, right? Here's the auditory nerve fiber. And it's hooked up to an inner hair cell. And it's sending messages. What are the messages? Neurotransmitter.

When the wave form goes like this, the auditory nerve fiber is responding. Ah, it's getting lots of neurotransmitter. Well, that was when the stereocilia were bent one direction.

Ions flowed in. The inner hair cell was depolarized. It released lots of neurotransmitter.

Let's go a little bit longer in time to this bottom part of the phase curve. The stereocilia were bent the opposite direction. The ion channels closed off. The inner hair cell went back to its rest-- minus 80 millivolts, let's say. And it said, I'm not excited anymore. I'm going to shut off the flow of neurotransmitter. The auditory nerve fiber goes, oh, we're quiet. We don't need to respond.

Go back to the other direction, then the stereocilia back the other way. Ah, I'm depolarized. I'm going to go to minus 30 millivolts. Ah, well, let's release neurotransmitter. Oh, wow, there's something going on. I'm going to fire. I'm going to fire all these action potentials.

It's going back and forth, back and forth. At some point though, this is going back and forth so fast that this just gets to be a blur. There is a sound there. It's depolarizing the hair cell. But it can't do this push pull kind of thing. It's not fast enough.

Even though there's a nice synaptic ribbon there to coordinate the release of the vesicles, it gets overwhelmed. Remember, at 1,000 hertz, this is going back and forth in 1 millisecond. And 5,000 hertz, it's going back and forth five times in 1 millisecond. That's pretty fast. And it gets overwhelmed.

There's a response. There's more action potentials with the stimulus than without. But they're no longer synchronized. It gets overwhelmed. And phase-lacking goes away.

We can distinguish 5,000 from 6,000 hertz very nicely when we listen. We're not using this code, because there's no temporal synchrony in the auditory nerve at

very high frequencies. This is a kind of an interesting code for sound frequency, because the timing is going to be different for different frequencies.

Imagine at low frequencies-- and imagine just for the sake of argument-- that the auditory nerve fiber is going to respond on every single stimulus peak. Let's say this is 1,000 hertz. And now let's say we dial in 2,000 hertz, which is going to end up going twice as fast. I'm not a very good artists here. But you can imagine that the firing is going to be twice as often, if for the sake of argument we're firing in every stimulus frequency, which may not happen.

But this is kind of an interesting code. Because if you're sitting in the brain and you're getting firing very far apart, you're going to say, OK, that's a low frequency. But if you're getting firing very close together, you're going to say, oh, that's a higher frequency.

So is there some little detector in the brain that's detecting these intervals? How fast the firing? Well, we don't know that. But we certainly know that a code is available in the auditory nerve at low frequencies, but not at high frequencies like 5 kilohertz.

So what's the evidence that we're using one code or the other? Clearly, the place code has to provide us with frequency information at a higher frequency. There is no temporal code at those high frequencies.

Down low, which code do we use? Well, we probably use both. That's another way of saying, I'm not really sure.

But let me give you some data from musical intervals that might suggest that this time code is used. What are the data? We have to talk a little bit about perception of musical intervals.

And we might as well start out with the most important musical interval, which is the octave. Does everybody know what an octave is? Yeah, what it is an octave? I can't explain it, but I know it.

What about on a piano? You go down and hit middle C, where is the octave?

**AUDIENCE:** The next C.

**PROFESSOR:** The next C, right. You've even called it the same letter, because it sounds so similar. But in precise physical terms, an octave is a doubling of frequency.

Whatever frequency middle C was, if you double that frequency, you get an octave above middle C. So we have some data here for two intervals, one 440 hertz-- two frequencies, one 440 hertz and another an octave above, 880. So double it.

And why did we pick 440 hertz? So that corresponds to a note-- just-- yeah, right, it corresponds to A. And I was trying to think if the A is below or above middle C. I think it's above. So what's important-- you guys knew that right away. What's important about that?

**AUDIENCE:** Orchestras tune to it.

**PROFESSOR:** Orchestras tune to it. So can you give it to me? OK, I'll give it to you. [WHISTLES] Sorry, that's A 440.

And so here's A 440 on the violin. [PLUCKS A NOTE] OK, now, how do I know that? Because orchestras tune to it. So for about 20 years, I sat in an orchestra. And the first thing you did-- [LAUGHS] OK, tune, you guys.

And what instrument gives the tuning note? If you're in junior high, it's this little electronic thing. But if you're in the BSO, what instrument gives the tuning note?

**AUDIENCE:** The violin.

**PROFESSOR:** No, violins go out of tune like crazy.

**AUDIENCE:** Oboe.

**PROFESSOR:** Oboe, right, because the oboe's a very stable instrument. And if the barometric pressure goes up and the humidity goes down, the oboe's still going to give you A 440. So the A 440 is a very important musical note.

And all these instruments, of course, have a whole bunch of harmonics. This string is vibrating in a whole bunch of different modes. But the fundamental, the length where the whole string vibrates is A 440. OK, so here's A 440 or approximately.

Now, an octave above that is a very nice sound. It's another A. That's A 880, the fundamental. And if I sound them together, they sound very beautiful.

And in any musical culture, an octave is a very predominant interval, because it sounds so wonderful to your ear. And violinists, I can tell you from experience, practice a lot of time trying to tune their octaves perfectly. And if you've ever listened to a professional go like this, and every time they go up and down, the octave is just beautiful.

But if you've been to middle school or elementary school, it's a little different. Because sometimes when those students play an octave, it doesn't really hit to be exactly an octave. And now I'm going to give you a demonstration that's 440 and not quite 880. OK And it's not going to sound exactly the same. So here it is.

And that's an interval I've listened to many times. But it's not a desired interval. It's a very dissonant interval.

And what is terribly displeasing about something that's not quite an octave versus an octave. That is a question that the place code has a lot of problems with. Because, for example, along the cochlea there is a place-- it's quite near the apex-- for the 440. And then if you go more basally, there's another place for the 880. And there's a place for the 879 and 878. And those would be very dissonant.

But there's no reason that those two things have any links to one another in the place code. There's a place for 1,000 and a place for 2,000. Why do they sound so wonderful together?

The timing code though has an answer for that. And here is some data to show you why those two intervals [? meld ?] very good together. If you look at the spike pattern, in response to either one of these frequencies, and compute what are called the intervals between the spike-- so-called interspike intervals-- every time

you get a spike, you start your clock ticking. And that interval is timed until the next spike fires. That's an interspike interval.

And obviously, if this is phase-locked, these intervals are going to have a close relationship to the stimulus period. So here's a spike and here's an approximately two-cycle interspike interval. Here's a short interval, but it's one complete phase. Here's a long interval, but it's now three complete phases. Here's another three-phase interval. Here's a one-phase interval.

You could make a very nice plot of the interspike interval in milliseconds, the time between the spikes. And these are the number of occurrences on the y-axis. So for 440 hertz it's the dashed curve here. And you get a big peak here that's a multiple of the period. So at 440 hertz, the sound wave form is taking about 2 and 1/2 milliseconds to go through one complete cycle.

And these intervals would be firing on successive periods, which, obviously, the nerve fiber can do. But sometimes they take a break, and they fire only every other period. And that's a double of this interval. And so you have a lot of firing at about 4 and 1/2 milliseconds, a lot of firing at about 7 milliseconds, and so on and so forth. So this is an interspike interval from auditory nerve firing in response to this low frequency.

Now, let's double the frequency. Now, the sound wave form is going back and forth twice as fast. And you have-- no surprise-- firing, in some cases, twice as short intervals. So here's an interval for the 880 hertz that's about 1 and 1/2 milliseconds.

But here we have a firing pattern that's exactly-- within the limits of experimental error-- exactly the same as for the 440 hertz. Here we have an interval that's representative of skipping a stimulus waveform or two stimulus waveforms for 880. But here we have a peak at exactly the same as 440 hertz, because these intervals are lining up every other one for the presentation of the octave.

When you put those two sounds together, you're going to get the combination pattern. And many of the intervals are going to be precisely on one another. And

that is a very pleasing sensation for you to listen to.

If you look at other very common musical intervals, like the fifth or the fourth, you will have many overlapping periods of interspike intervals in auditory nerve firing. And those are very common musical intervals. If you look at dissonant interval, like 440 and 870, there will be no overlap amongst those two frequencies in the auditory nerve firing.

Now, let's go back to psychophysics and give you one more interesting piece of the puzzle here for why temporal codes might be important is active matches become more difficult-- and actually impossible-- above 5 kilohertz. OK, well, how does that fit in? Well, we just said that phase-locking-- because the hair cell and auditory nerve can't keep up with one another, the phase-locking diminishes for these high frequencies. And there it becomes impossible to match because you don't have this timing code in the auditory nerve.

So most of musical sounds are confined to the spectrum below 3 kilohertz. If you look even at the upper limit of the piano keyboard, you're sort of right at 3 kilohertz or so. And that's probably the reason. It's a very likely reason.

Now, we have a research paper for today. And I'll give you the bottom line or the take home message for that. And this is an interesting neurophysiological study based on a psychophysical phenomenon called the octave enlargement effect.

So I wasn't quite truthful when I told you that octaves are the most perfect interval to listen to, because it turns out if you have people-- if you give people a low tone and give them an oscillator. And say, dial in an octave above that, they'll dial it in and say, ah, that sounds so great.

But if you look really carefully, it's not exactly an octave. It's a small deviation. What they dial in-- especially at high frequencies-- remember, you can't do this at really high frequencies, but at 2,500 or 1,500 hertz, toward the high end of where you can do it-- they dial in actually a little bit more than an octave. And they say, ah, that sounds great.

But it's not exactly an octave. So this paper looked and said, what about auditory nerve fiber firing can explain this octave enlargement effect? The fact that people dial in a little bit more than an octave for the upper tone.

So these are the psychophysical measurements. And they just give you previous studies. What they did was they recorded from the auditory nerve. And they looked at interspike interval histograms, like we've just been talking about.

And they saw something really interesting at the very high frequencies. So they're going to especially concentrate in here. They found that the very first interval didn't match exactly what was predicted.

So this is a stimulus of 1,750 hertz, so toward the right end of the graph here. And where you'd predict the intervals to happen is shown by the vertical dashed lines. And they didn't fire right on those predictions, except for the very shortest intervals. And they said, what's going on here?

Well, when you get to very high frequencies, what are we talking about for these intervals? Even at 1,000 hertz, what's the time scale here? This is 1 millisecond.

What problems do you get when you ask a nerve fiber to fire and then you ask it to fire again a millisecond later? That's a very brief interval. Anybody know? What is the limit of firing? Can nerve fibers fire closely spaced action potentials, you know, less than a millisecond? What's the problem that they have?

**AUDIENCE:** Is it that they are polarized?

**PROFESSOR:** They're hyperpolarized, right. What else?

**AUDIENCE:** Because of the refractory period.

**PROFESSOR:** That's right. There's something called the refractory period, which can cause them to hyperpolarize. So what's happening-- if this were a nerve cell membrane, is what happens is sodium channels can open up to allow sodium to come in and depolarize the neuron and fire an action potential. And then those channels are

turned off. And potassium channels open up and allow potassium to go back and even hyperpolarize.

But these channels take a little bit of time to recover. It takes a little bit of time for the sodium to turn off and get ready to fire another action potential. It takes a lot longer time for the potassium channel to close and get ready to fire another action potential. And the refractory period is the time it takes for everything to recover fully so we're ready to fire again.

And in the limit, that's supposed to be about 1 millisecond. That's the absolute real refractory period. So when I was drawing here things a millisecond and less, I wasn't really being truthful.

There's something called the relative refractory period, which is a couple of milliseconds. And the nerve fiber can respond, but it's not going to respond quite as quickly as before. All these channels aren't completely reset. It's going to take a little bit longer time to respond. That's what's going on in this very first peak.

Remember, this first peak indicates firing at successive cycles of the sound waveform at 1,750 hertz. It's a very brief interval. And what happened is you fired. And then the next action potential you fired, but you delayed a little bit. So the interval is a little bit longer. That pushed this one up to the next peak, so that this interval was actually too short.

What the brain is getting is an interval that's a little bit too short. When it hears that, it says, I want to recreate the higher frequency to be an interval that's a little too short, I'm going to dial in a little bit too high in frequency. OK, because that frequency sounds better with this too short of an interval.

So go ahead and read that paper. It's a very interesting study of how neuronal firing can give you a psychophysical phenomena. That's quite interesting. And it happens especially at the high frequencies. Octave matches the low frequencies are sort of as you would predict as is auditory nerve firing.

OK, any questions? If not, we'll meet back on Wednesday. And we'll talk about the

cochlear nucleus, which is the beginning of the central auditory pathway.