PROFESSOR: OK. I guess we'll get started. So last time we were talking about the auditory midbrain, where the big structure is the inferior colliculus. And we talked about its inputs coming from the superior olivary nuclei, like the lateral superior olive, which has neuron sensitive to enter oral level differences, and another input being the MSO, where the neurons insensitive to binaural stimuli with interaural time differences. And those inputs converge on the inferior colliculus.

And in today's lecture, we have a little slide of the inferior colliculus. And the big part of the inferior colliculus, diagrammed here is this central nucleus, the ICC. And that structure has been well studied, part because it's the big part. It's easy to record from, it's deep in the brain. It has a strong tonotopic organization. We talked last time about neurons there having time intensity trading because they get in some cases, input from both LSO and MSO.

We talked about coding for stimulating head precedence like effects in the inferior colliculus, and we talked a little bit about sound localization in there in the mammal, at least for the parts of the inferior colliculus that had been studied, there being no really good or obvious map of sound location to place in the colliculus, and we also talked about some of the other edge regions of the colliculus like the external nucleus of the colliculus that has been explored very nicely in the bar now.

And we talked about its projections up to the optic tectum, or the tectum in the bar now, which is the analog of the superior colliculus in the mammal. And there's this beautiful mapping of auditory space in the barn owl at that position in the brain. Now, such mappings are not found in general in the auditory system of mammals, with the one exception being in the deep layers of the superior colliculus.

There are neurons that are responsive to auditory, as well as visual stimuli. Of course, we think of the superior colliculus as being a visual nucleus. And in those deeper layers, there is a mapping of auditory space that's in line with the mapping for visual space there. So it's as if a mapping of space can be created in the mammal, but it just doesn't appear, at least in the main nuclei of the auditory pathway.

The superior colliculus is really not an auditory nucleus. OK. So any questions from last time? So today, we're going to be talking about mainly the auditory cortex and how it has a number of fields in the cortex. Usually we talk about fields rather than nuclei, and these fields are tonotopically organized-- at least, several of them are. And some of them aren't, but we'll be emphasizing the tonotopic ones.

There are some very interesting experiments where investigators have explored how changes in hearing can affect the tonotopic mappings in these fields so that they can be plastic. They can be molded or shaped depending on experience. So we'll talk about that plasticity. Because there are many fields, the obvious question is, well, what does one field do in your sense of hearing, and what does another field do.

And we're generally not able to answer those questions. I like to think that there are Nobel Prizes waiting to be earned in this area as that function is worked out for each of the fields. But it's clear that there's a big role of one field, which we call A1, in sound localization. And we'll review that line of investigation. That if A1 is lost, an experimental animal's ability to localize sounds is completely gone.

OK? So we'll talk about the evidence, how we test animals for sound localization. Then we'll talk about where A1 is in the human auditory cortex. And we'll talk finally about some very interesting imaging studies that have shown that there is a center that lights up in imaging studies when you present sounds with salient pitch. That is, very strong sensations of pitch, and it's an area near A1, but a little bit beyond it.

So we'll talk about those imaging studies. And that relates to the paper that we have for assigned reading for today's lecture. So that's the roadmap for today. So the first

slide is a very complicated slide that as I said, has the inferior colliculus on it. And in the second row, has the auditory thalamus on it. So the auditory thalamus is the medial geniculate. And I'm not going to say very much about the medial geniculate. Probably Peter Schiller talked a lot more about the lateral geniculate in the visual part of the course.

I want to just say that there is a large part of the medial geniculate called the ventral division, indicated here. There's a couple little symbols that say v. That stands for ventral division. And there is a large part of the medial geniculate that is tonotopically organized. And so it receives a lot of input from the central nucleus of the colliculus, which is tonotopic. It is tonotopic.

And it, in turn, projects 2 auditory cortical fields that are listed up there, which we'll go over in detail in just a minute. And all of those listed cortical fields, with the exception of the one that says t, all of those are tonotopically organized. So to some people than, there is a part of this auditory pathway at these higher levels that's called the tonotopic system. And that has a bulk of the nuclei.

There are some other systems. And at least in this author's idea, they're labeled as diffuse and multisensory systems. We're not going to talk much about them. But they start out in other parts of the colliculus, like the dorsal part of the colliculus, and the external nucleus of the colliculus. And they pass through other divisions of the thalamus. And they either go to other auditory cortical fields, like A2 here, or they project diffusely to a whole bunch of places in auditory cortex.

So we're not going to emphasize them very much at all. Just stick with the tonotopic system for now. So in the auditory cortex, there's been a lot of research done on the cortex of the cat. Cat has been a standard model for auditory cortex work since the beginning, since the '60s. And so the fields are very well known in the cat. And they're indicated in the colored shading here. And this is showing you the cat's brain from looking at its left side.

So you see that the colored areas, the auditory areas, are in the part of the brain called the temporal cortex, the side of the cortex. So way up here would be the

frontal part of the cortex, the frontal lobes here. Way back here, you're familiar with the occipital cortex, which would be the visual areas. And these areas are on the side, our so-called temporal cortex. Those are the auditory areas.

And the cat is a nice model for cortex. Because much of the big auditory cortical field, A1 here, is on the surface. It's on a big gyrus that's accessible. You can get to it, you can record from it. You can make lesions in it. A little bit of A1 goes down into a sulcus. And those sulci are labeled by the lines. This one is labeled ps, the posterior ectosylvian sulcus.

And this is the cortex with the cortex map pulled apart. So you can see the parts of the cortex that's down in the sulcus. And those are the dark pink here. And in the front part, rostral part of the auditory cortex, there's another anterior ectosylvian sulcus. And that's shown by the dark pink here. So much of A1 is on the surface, but a little bit of it dives down into a sulcus.

Now, when you go with microelectrodes and record from the auditory cortex, you take off the skull, you take off the dura, and you see the cortex. And you can take your microelectrode and go right into the cortex. And of course, the cortex has layers. So you'd be going, starting in layer 1, and going down. How many layers are there in cortex? There's 6, right?

OK. And you find that if you go down through all those layers and make recordings in the different layers, all of the neurons in one cortical penetration with your electrode have the same CF. So of course, we want to measure the characteristic frequency from the tuning curve, just like we do in other places. And if you do that, in these recordings you find that they all have the same CF.

And that's like going down through a cortical column. So columnar organization is very important in cortex. And if we draw it from the side, this would be the surface. And you'd have these 6 layers usually labeled by Roman numerals. And your electrode would be coming down here and sampling from the various layers, single neurons at a time.

Of course, you'd have to maybe record from layer 1 first, then layer 2, and then so on, and so forth. And as long as you stay within a column of cortex, the neurons have the same characteristic frequency. So looking at it from the surface here then, one column would be like a dot. You'd be looking down at the very top of the column from the capital down to the base of the column, if it were an architectural column.

When you do that, you can make a CF mapping for these tonotopic fields. And you find that if you're way over here, rostral in A1, you get very high CFs. And so high CFs for a cat would be like 40 kilohertz, maybe even 50 kilohertz. An octave beyond the upper limit of human hearing. But if you move your electrode, more caudally, the CFs get lower, and lower, and lower until you reach the posterior edge of A1.

And this little word, "low" means, that's where the lowest CFs in A1 would be found. And the low CFs would be like 0.1 kilohertz, 100 Hertz. If you went perpendicular to that so-called tonotopic axis, you can have recordings that go one here, one there, one there, one there. If you're moving up from ventral to dorsal, then you find they all have the same CFs. So that's an iso CF.

Lamina, if you will. And it starts right here. This line right here is an iso CF lamina. This would be the highest CFs. And then as you go caudally the CFs get lower, and lower, and lower. Then, when you keep moving your electrode, more and more caudally, you find further CFs. And then the progression changes. The CFs start to become higher, and higher, and higher.

That's the signal that you've entered another auditory cortical field right behind A1. And it gets its name from being behind it, behind an anatomical terminology is posterior. So this field is p, or sometimes called PAF, posterior auditory field. And if you keep going, more and more caudally. In this case, the thing takes a turn. So you go, caudally and ventrally, the CFs then start to get higher, until you approach the boundary of p with the next auditory cortical field, which is VP.

Ventral posterior auditory field. Then the CFs go from high to low. Same thing happens where A1 meets the anterior auditory field, indicated by a here. Those two fields share a high frequency, high CF boundary. And then, as you go more rostral,

the CFs get lower, and lower, and lower. Sometimes this organization, at the edges of the fields, is called, mirror image tonotopy in the auditory system.

And where else have you guys seen this? Did you go over this in the vision part of the course where you have the mapping of the retina on V1 in the occipital cortex, right? And it has a retina topic mapping. And that has some image where the nasal part of the retina is over here, and the temporal part of the retina's over here. And where V1 abuts V2, you also have a retina topic mapping.

But it's a mirror image reversal of V1. So you should have gone over that in the visual part of the course. You find such mirror image flips in the somatosensory cortex. In the somatosensory cortex, there's a mapping of the body surface. If you touch here on the body surface, a certain part of the somatosensory cortex responds. If you touch up here, a different part responds.

And there's a mapping of the body surface onto the surface of the cortex. And where S1 meets S2, there's also a mapping. But it's a mirror image. OK. So this is not a surprise in terms of general cortical organization. And it's also not a surprise in the visual periphery. You have the retina topic mapping. In the auditory periphery, you have the CF mapping. And you have this nice CF mapping along the cortex.

OK. And in the cat, you have these four tonotopically organized system fields. And you have several fields; A2, DPV, and T, where the tonotopy is either nonexistent, or much less obvious. And there are some challenges here to exploring responses in these other areas. For example, tuning curves in A2 can sometimes look like this. And it's very difficult then, to assign a CF to such a broad bowl shaped tuning curve.

You could do it, if you really were pressed. But there's not much difference between that frequency and that frequency. So this could be an octave or more difference. And so it's hard to assign a CF to some of the neurons in these other fields. That's not true in A1. In the tonotopically organize fields in general, there's very precise, sharp frequency tuning. OK. So the tuning curves, there's a little table here.

They're usually sharp in these tonotopically organized fields. Yes, there's tonotopic

organization. The latency is short. The response is brisk or robust compared to response that might be called poor or insecure. That just simply means, every time you turn on a sound, and people remember, typically tend to measure histograms in response to hundreds of sound bursts every time there's a response.

But here, there might be a response to the first tone burst in a train. And then the neuron might shut down and not respond anymore. So these diffuse areas are hard to investigate. Most of these cortex recordings I'm talking about have been done in anesthetized animals. And so there's always a question of, how much has anesthesia changed the response patterns? Of course, anesthesia has a big effect on these higher levels of the nervous system.

So it's not always clear how much of these properties or change in these properties has been due to anesthesia. So back to the somatosensory cortex. Here's a mapping of the somatosensory cortex. And this is the surface of the body map onto the surface of the cortex that I was referring to earlier. So if you stimulate the surface of the body over here, the tips of the toes, you get a response here in the somatosensory cortex.

If you stimulate the fingers, you get a response here in the somatosensory cortex. You stimulate the facial region, you get a response over here. And this distorted mapping or caricature of the body's surface, sometimes called a homunculus. And it shows you that a lot of cortex is devoted to certain important parts of your body, like the face. And less important parts, like the trunk of your body, receives much smaller representation in the cortex.

In the auditory cortex, if you draw the mapping-- so here's a mapping of A1. In this case, it's from the guinea pig cortex. And in this case, the mapping is plotted with the CF on the y-axis, and the distance along the cortex on the x-axis. And these are the data. Each dot indicates a recording site from a single column in the cortex. And if you could read these numbers, they would indicate the CFs. And these lines are the ISO CF contours.

And this CF axis is plotted along this distance here. So this distance is going like

this, and this distance is going along like this. And there's a very nice, almost linear relationship. And it shows you something quite different from the somatosensory mapping, which is that there aren't any really important frequencies. They all have about the same representation in the cortex.

It's a pretty boring, straight line, if you will. It's not as interesting as the homunculus in the somatosensory cortex. So this is true in the general mammal. Next time, when we talk about the auditory cortex in the echo locating bat, we'll have quite a different finding. There are some very important frequencies in certain types of bats that relate to the echo locating signal that they emit, and the echo that comes back to them so that they can find targets, even in the dark.

But most of general mammals do not echolocate, of course. And so their mapping of the sensory cortex is pretty linear and boring, if you will. So today's reading comes from what motivates the next experiment that I'll show. For a long time, these mappings were thought to be laid down at birth and not changeable. So they were just immutable. But some very interesting experiments by Dexter Irvine and Don Robertson in the 1980s showed that was not true.

And they were not the pioneers in showing that cortex can change as a result of experiment. Rather, some people who were working, especially in the somatosensory cortex, were the first. And so I didn't have a reading for today. So right before I came over, I pulled up the Wikipedia entry on the Silver Spring monkeys. So has anybody heard of the Silver Spring monkeys?

So Silver Spring monkeys were in the news in the 1980s. They got their name from the Institute of Behavioral Research in Silver Springs, Maryland-- Silver Spring, Maryland. And from 1981 to 1991, they became what one writer called, the most famous lab animals in history, as a result of battle between animal researchers, animal advocates, politicians, and courts. So there was a researcher whose name was Edward Taub. And he was experimenting on the somatosensory sensory cortex.

And he was taking the monkeys, and he was denervating the sensory input from

certain parts of their limbs. So for example, he would cut the nerves that carried information from the middle finger of the monkeys. And he was studying to see if the somatosensory cortex remapped, and he was finding small effects. But in May, 1981, Alex Pacheco, from the animal rights group PETA, began working undercover in his lab, alerted the police to what PETA viewed as unacceptable living conditions for the monkeys.

And there was a long battle. Initially, the researcher was convicted of animal cruelty, and these charges were subsequently overturned. But anyway, the monkeys were held in limbo for in some cases, many years because his research was put on hold. During the subsequent experiments on the monkeys after the court battles were all done, it was discovered that significant cortical remapping had occurred.

This is evidence of the brain's plasticity, and it helped to overturn the widely held view that the old adult brain cannot reorganize itself in response to its environment. So the analogous experiments in the auditory system have been done in small animals. And maybe it's a result of much decline and use of primates in research. There was hardly any auditory work done on primates these days. This reorganization work is done in the guinea pigs. And the experiments are done like this.

There's a peripheral lesion made in the cochlear. In the Guinea pig cochlear, it's very easy to make a little hole in the middle layer, and look down on the cochlear. And the cochlear's a bony structure. The most accessible part is the basal turn of the cochlear. And you can go right through the round window and see the basilar membrane, and the hair cells. And you can make a little tiny pinpoint opening in the organ of corti with a fine metal pick, and create a substantial hearing loss in one little place of the cochlear.

So here is indicated a graph of the compound action potential threshold. This is a response from the auditory nerve. Action potential, obviously, is an impulse from single auditory nerve fibers. Compound means it's a recording from many, many, if not all of the auditory nerve in response to a tone burst of the different frequencies.

And if you make a small lesion at the basal turn, remember the frequency organization of the cochlear is such that the basal turn processes, high frequencies. And so instead of the normal curve in the lesion animal, you have a big increase in threshold, maybe 60, or 70, or 80 dB in the lesion case. And that lesion goes from about 10 kilohertz to about 20 kilohertz.

And in other parts of the cochlear, the hearing is normal. So this is a peripheral hearing loss. And now, we're going to then look in the cortex and see if the tonotopy of the auditory cortex is the same or different. And obviously, my big build up here, that there's plasticity of tonotopy, is found. This is the normal mapping that we first saw. This is the mapping in the lesioned animal.

And in this case, it's a mapping where each of these dots, of course, is a recording site. These are very high CFs. You march along here, and the CF gets a little lower. But there's a big region of about 20 kilohertz-- big, long distance in the cortex-- when all you get is CFs of 20 kilohertz. Notice that that is right at the edge of the lesion. In the lesion, between 20 and 10 kilohertz, you don't get any response.

Well, there's a huge hearing loss there. It's no surprise that there's no response to those frequencies. The auditory periphery is not sending you any messages, or sending you very few messages about those frequencies. Then you jump a little bit in distance. And the CF jumped from 20 down to 10 kilohertz. And there's a long region of cortex at which the CFs are all 10 kilohertz.

And notice that 10 kilohertz is a very important frequency in the audiogram of this animal, and that it's right at the edge, the low frequency edge of the hearing loss. After that extensive region then, you pick up your normal tonotopy of auditory cortex. So this is clearly a massive reorganization compared to the normal of this lesioned animal's cortical mapping. So couple comments about this. If you do the mapping right after the lesion, this reorganization is not found.

So it takes some time. In the case of the auditory system, it takes more than three weeks to see the reorganization. What's the mechanism for this reorganization? Well, we have input coming up to here from the thalamus, where the inputs from the

20 kilohertz place of the thalamus, did they come up here? And did they grow into a large part of the cortex where they weren't present before?

And did they do that growth because that part of the cortex had gone silent because of the hearing loss? So one mechanism could be growth, if you will, sideways growth of axons. Another mechanism could be the inputs from the thalamus are coming up here. And even in the normal case, they don't just go to one place, but they have all sorts of side branches. And these side branches in the normal case are inhibited, or they're not strong to begin with.

Because there's a lot of other stuff coming in here, and the other stuff is saying, you guys be quiet. I'm the main input. And maybe after the hearing loss, that other stuff is silenced. And these previously weak inputs become stronger. So you could say then that another mechanism is strengthening of preexisting inputs that are weak before the perturbation. And I wrote mechanism, ? because we don't know what the mechanism is.

And it may be both. They're not mutually exclusive. And we're thinking here, cortex. This is a lecture on the cortex. But I should bring up, where is the locus of this plasticity? So in the auditory system, we have a rich array of nuclei. We have the cochlear nucleus, we have the superior olivary complex, we have the inferior colliculus. We've just learned a little about the auditory thalamus.

Well, could one of those lower level nuclei have reorganized, and then just passively spit their reorganized input up to the cortex? And the answer is yes, and that actually has been looked at. So there's no reorganization in the cochlear nucleus. The part of the cochlear nucleus that processes 20 kilohertz is supposed to be normal. In between 20 and 10 kilohertz, you have a normal region. But it's completely silent in these types of hearing loss.

You have a small amount of reorganization in the inferior colliculus. But it's not as big as what we see in the cortex. And maybe in the medial geniculate body, we have a larger reorganization, but probably not quite as much as you do in the auditory cortex. So there may be a little bit of reorganization, an IC, a little bit in

MGB, and then further reorganization in auditory cortex. So that has been looked at. And the answer is at all of these higher levels, there's reorganization.

But there's maybe more as you go higher in the pathway. Now, who cares? We usually don't have hearing losses, right? Well I should remind you that we had a big lecture on hearing loss. Then let me show you one of the slides from it. So we had this slide under the lecture on hearing loss. Well, I've given it a new title now because I was thinking about getting old this week.

And the type of hearing loss that you have when you get old is called presbycusis. Presbycusis is the age related loss of hearing, especially at high frequencies. And almost all of us are going to go through presbycusis. This is a normal audiogram or threshold of hearing curve when you're young, and this is one when you get older. And invariably, we lose our high frequency hearing.

The causes are not known. But clearly it takes place in the periphery, which is very similar to the lesion study we just went over. It's a peripheral hearing loss, what happens to your central pathway. It probably results in cortical reorganization in the human. And again, I can't spell. There's a missing t here. So we probably all-- and we have plenty of time. This hearing loss doesn't happen in days. It happens over the course of our lifetime.

There's plenty of time to reorganize the cortex. Now, let me go back a little bit here and ask an interesting question which has some negative answers, and some positive answers. People have said, well, wow. This animal-- you have a lot of cortex. Half a millimeter devoted to CFs of 20 kilohertz. And you have a lot of cortex devoted to CF of 10 kilohertz. Does this animal do something a lot better at those frequencies than this normal animal which just has a little bitty part of cortex devoted to 20 kilohertz, and a little bit devoted to 10 kilohertz.

So it's not clear what the answer is to that question yet. And people have speculated in the normal-- if you train a normal person or a normal animal to do a task at 10 kilohertz. So they're listing to 10 kilohertz over and over again. It's clearly a training effect for many tasks. You get better with training. Does that mean in the normal

case that we enlarge the 10 kilohertz part of our cortex?

Well, it's not known. Does that means this big area of 10 kilohertz and 20 kilohertz in the lesioned animals cortex enable that animal to do something better? That's not clear. There's evidence for both answers to those questions. Let's move on to some other properties that have been observed in auditory cortex. And we've been stuck in this mode and this course on frequency organization.

It's clearly a very strong component of central auditory nuclei. And here are some tuning curves from auditory cortex neurons. You have the x-axis being frequency, and the y-axis being sound pressure level, response being inside the tuning curve. In this case, they've embellished the tuning curve a little bit by plotting the biggest, best response in real black shading. And these are the tuning curves we've been talking about in the course all along. You might see goes from the auditory nerve from cochlear nucleus and lower levels.

In the auditory cortex, you see some of those kinds of tuning curves. But you also see from different neurons, other types of tuning curves. And these are seen to a greater extent in the auditory cortex, although you do see some, to a certain extent, in the colliculus, and to a certain extent, in the thalamus. So you see them in a numerically greater extent in the cortex. And here is a tuning curve that sure has a best frequency.

But at that best frequency, if you get it higher and higher in level, the neuron actually stops responding at a pretty moderate, and certainly at a high level. And if you study its response inside the black shading, you can say, well, that neuron has the best frequency, or characteristic frequency. It also has a characteristic level. It likes it a lot when the tone frequency is 10 kilohertz, and the tone level is 40 dB. That's its maximal response.

That's the one it prefers, if you will. Here's a different one. A best level is somewhere in the middle here. Here's a best level here in a very narrow response area. So many of the neurons in the auditory cortex has to have these not only characteristic frequencies, but also best levels. If you, at their CF, raise the sound

level, you get this type of pattern. These are rate level functions. So this is the firing rate, and this is the tone level or tone intensity.

And this is from a number of neurons. But just concentrate on this one. The firing rate goes up with level, reaches a maximum, and then it declines. And at the highest level, it doesn't respond anymore. So we don't know what this means. It means somehow that these neurons could tell the animal or the person, well, the sound level is x dB. They can tell you the sound frequency is a certain number of kilohertz in the sound level if they're responding maximally, is a certain SPO.

It's clearly very different. Now, I brought that up because people have looked at that in terms of coding for preferred areas of space where the neuron's response is maximum. And these are some data from Clarey et al on azimuth level response areas for cortical neurons. So what does that mean? Well, they're recording from a single neuron in the cortex. And they have the animal-- in this it's, a cat-- in an anechoic room.

And they move the sound source around in azimuth. We've talked about this with experiments in the [INAUDIBLE]. And they study the neuron's response as a function of azimuth. In this particular study, they also varied the sound level. And so that's what the y-axis is here. This is the sound level axis. And what we just said is that many neurons in cortex have a preferred sound level.

They start to increase their firing. And that's what's meant by this shading here, the dark of the shading, the higher the firing rate. And then at high sound levels, the firing trails off, or it goes down to zero. This is a level function for these different neurons. These are recordings now from the auditory cortex within a single column. So shown here is the electrode penetration going from unit one down two unit 10.

And those are indicated here, unit 1, down to unit 10. And the cortical layers are indicated by the Roman numerals 1 through 6 here. And they're recording these 10 different neurons from a given cortical column. And what's impressive about these studies is that the azimuth of each of these neurons, where it prefers in space, it is it a certain azimuth, 45 to 90 degrees. And that's as if you were recording from the left

auditory cortex.

45 degrees would be over here, to 90 degrees would be straight over here. So these neurons are going to respond when the speaker is in a position over here. And they're going to respond when there's a moderate sound level. Not the lowest level, and not the highest level. So they prefer a certain sound level. And within a given column, almost all of the neurons have similar types of azimuth level functions.

And remember, before we said that these all have about the same CF. So a second thing that is common to units in a given column and auditory cortex is their azimuth level response areas. And that's shown very nicely here. These type of data suggests that maybe these neurons play some kind of role in telling us where a sound is coming from. Like without this column, maybe we really wouldn't know that sound sources were located at 45 to 90 degrees over there on the contralateral hemifield.

So there's been a lot of work on auditory cortex and sound localization. And I want to get into it here. So how do experimenters test behaviorally for how an animal can localize sound? Well, this is the formal way to do it. Here's an experimental animal. It's going to a speaker that emitted a sound. Before, it had been sitting in this central position in the testing cage here.

Waiting maybe cued by a light that's saying, the trial's about to start. And the animal then listens. And this speaker up here near b didn't emit the sound. But the speaker down here did emit the sound. And then the animal is trained to go to the speaker that it had heard emit the sound. If it does that correctly, there's a little food reward area down below the speaker, and the animal gets a food reward.

And the animals are food deprived, so they're motivated to do this task. In this case, the speakers can be moved. So you have removable speakers. Or you can have an array of speakers, as indicated here. And the animal has to choose which of the several speakers emitted the sound, and go to the correct one to get the food reward. And this has been done with a variety of experimental animals. In this case,

it looks like a cat. But they've also tested rats and monkeys.

Now, there's a couple things you have to worry about here. You have to worry if your sound is on a long time that the animal isn't cheating. And one way of cheating would be for the animal to sit still here, and listen to the sound, and then move a little bit, maybe just by bending over, saying, OK. If I moved over here, did the sound get louder? Oh, that means the sound is coming from this side.

So generally, these tests are using pretty short stimulae-- 50, 100 milliseconds. And during that period, the animal doesn't have a chance to move and sample the sound field. So what are the data in normal hearing animals? 75% correct at distinguishing which of the speakers have emitted the sound at 5 degrees. So in that case, it's when the movable speaker's just 5 degrees.

With a, in this case, 500 millisecond long spectrally complex stimuli. So animals are not as good at this sound localization task. In human performance, I think back a few lectures ago, we talked about the minimum audible angle in humans being a couple of degrees. Maybe 1 degree. In this case, the animal's going to 5 degrees. So it's not quite as good. The animal can still do the task surprisingly, even if it has just one functioning ear.

And how is that possible? Well, the animal's pinna. We talked about the external ear or pinna, providing you some nice spectral queues. And it looks like the spectrally complex stimuli are being used. So those spectra cues are available. But the minimum audible angle's more like 10 to 12 degrees, with just one ear with a good pinna on it. So the best performance is with two ears.

Now, why am I going over this paradigm? Well, people have then taken experimental animals and studied them after lesions. And we're talking about the auditory cortex. So let's look at the results of lesioning the auditory cortex on sound localization. So looks like in this study, they're using the array of loudspeakers. The lesion is located in the auditory cortex.

And this is the right side of the cortex. This is the occipital or back part of the cortex.

This is the front of the cortex. So the lesion is made on the right side. When a lesion is made on the right side of the auditory cortex, the animal has problems localizing sound in the opposite hemifield. So the lesion is made on the right side. The animal doesn't know if it's this speaker, this speaker, this speaker, or this speaker that's emitting the sound as a random performance on that side.

However, the animal can distinguish between the speakers in the hemifield ipsilateral to the lesion, which suggests that the intact auditory cortex is mediating that behavior. So there's a deficit contralaterally in the opposite hemifield. And as you can see from this lesion, which was located smack in the middle of A1, A1 lesions effectively knock out sound localization behavior.

So these early studies suggested that A1 is critically important, and is necessary for correct sound localization behavior. And in these early studies, the lesions were actually made surgically by taking out some cortex tissue. And they became, as time went on in the mid1980s, the lesion studies became more elegant in that before the lesion was made, frequency mapping of A1 was made.

And this frequency mapping is shown here. This is best frequency, or CF. Those terms are synonymous. And this is distance along the cortex. And not the entire auditory cortex, field A1, was removed, but just a particular part, just a little distance here. And it was known from the mapping what CF was affected by the lesion. And the other CFs were left intact.

You can test the animal for any frequency you want to. So you can test for frequencies in the intact area, or you can test for frequencies in the lesioned area. And it was shown very clearly then by plotting performance. This is a performance axis, where downward is very accurate. This must be the number of mistakes made. So 0 is no mistakes made. At the low frequencies where the cortex is intact. At the midfrequencies, where the cortex is lesioned, performance goes to chance.

And then at the high frequencies, again, where the cortex is intact, the performance gets very few errors. It's very good. So this elegant experiment shows then that sound localization proceeds by frequency independent channels. That is, the part of

A1 that's responsive to low frequencies is mediating low sound frequency localization. And the part that's responsive to high frequencies is mediating high frequency sound localization.

So a very beautiful demonstration. Now, these lesions were done with the techniques available at the time. And it's very simple to go in and destroy a part of cortex. And for one reason or another, people decided to revisit these lesion experiments, even though they were very convincing, done in many, many different species. They decided to revisit them with a completely different way of making a lesion. And this is the method of inactivation by cooling.

And maybe many of you have heard of this. This might be the auditory cortex, the surface. This is layer 1 and the different cortical layers. The way you can inactivate the cortex by cooling is by taking a piece of tubing that has cooling fluid, or if you will, coolant. And you can put this coolant with a pump through this tube, and you can lay the tube right on the surface of the cortex. And obviously, as the coolant comes through here, it's going to cool down first the top layer of the cortex, and then the lower layers.

And finally, all of the cortex. And you can assure yourself that this cooling has inactivated the cortex by doing things like recording evoked potentials. And what's elegant about the cooling experiments is that you can reverse them. So I don't know if this is a word, but instead of coolant, you can use a warmant, and restore this to body temperature. And responses come back. And what's very impressive is that these kinds of experiments can be done in animals that are actually doing a behavioral task, localizing sound.

And clearly, those experiments confirm these earlier lesion experiments that if you cool A1, you get a deficit for a localization ability in the contralateral hemifield. However, they have also come up with an interesting result in that if you cool some other fields, you also change sound localization behavior. Field PAF, when cooled, also interrupts sound localization behavior. So that's posterior to A1. And a small field that's not named, but is right on top of the anterior ectosylvian sulcus, the AES.

When that area is cooled, sound localization behavior's also disrupted.

So I used to be able to say, anyone who performs this critical function, sound localization, well, not so sure about it anymore. Because if you cool these other fields, you have a disruption of sound localization behavior. If you cool any of the other fields like A2, of VPAF, or most of the AAF, you don't get an interruption of the behavior. So what do we take home from that? Well, it seems like there are several fields that are important in sound localization behavior.

Looks like A1, p, or PAF as it's sometimes called, and region near the anterior ectosylvian sulcus are all important. And it's a little bit controversial, why the old lesions didn't actually show this. It seemed like in the old lesions, there were some studies which said, if you leave A1 intact, and you lesion all the other cortical fields, the animal can still do the task. That doesn't really fit with the cooling results, which has several fields that are important.

So don't let me leave you with the idea that what these fields do is only sound localization behavior. So A1 may be involved then hundreds of other tasks related to our sense of hearing. It's also involved in sound localization. So that's the right way to think about it, that these fields probably do many things. And I think if you got the gist of what Doctor Schiller talked about in vision, he's not a big fan of this little area of cortex does this little function. And over here, this little area does this function.

He's more of a believer in holistic cortex function where to do a task, you employ a lot of cortex; auditory cortex if you're doing an auditory task. And perhaps the more difficult a task is, the motor cortex you use. We'll see an evidence of that. And next time, when we talk about language processing in humans from imaging studies. So what else does cortex do besides sound localization?

Oh, I forgot to talk about auditory cortex in humans, how many tonotopic fields there are. So I brought this nice model of the primate brain. And where is auditory cortex in humans? So this is a slice of the brain, as if you were to cut it like this. And look at one slice. And this is the right and left temporal lobes. So in the primate, you have actually a separate lobe of the brain called the temporal lobe.

And in the temporal lobe, you have to-- the primate is a little bit more difficult to examine than the cat-- you have to pull down the sylvian fissure, which is between-- separates the temporal lobe from the parietal cortex up here, and look on the superior surface of the temporal lobe, and find the sight of A1. And on that superior temporal lobe surface, you have a little gyrus that was examined by an animus Heschl.

And Heschl's gyrus in humans is the site of A1. Some humans actually have 2 Heschl's, gyri, and they have their A1 either on one or both of the Heschl's gyri. Now, looking at it from the side view, in the temporal lobe, you have the 3 big gyri. Superior temporal gyrus, inferior temporal gyrus-- sorry, middle temporal gyrus, and inferior temporal gyrus. And so A1 is on the superior surface of the superior temporal gyrus on a little bitty gyrus called Heschls.

So I'm going to pass this model around. And A1 is indicated by a little piece of yellow tape there. You can take a look at it. And so that area, Heschl's gyrus, lights up very nicely in imaging studies when you present auditory stimuli. And so here's an imaging study where the imaging plane was parallel to the sylvian fissure or sylvian sulcus, and it's capturing just this superior temporal gyrus.

And the plane is looking down here. And on Heschl's gyrus, you see the left and the right in this case, and just the right in this case lighting up. And you can use, of course, different frequency sounds. And change the frequencies of those sounds in a progression from high frequencies to low frequencies. And you can draw the progressions that you see in imaging signals, and that's what's drawn with these arrows.

This is a MIT thesis by Tom Talavage. And he showed that there were at least 1, 2, 3, 4, 5 clear progressions of frequency sensitivity, as if you were progressing along tonotopically mapped auditory cortical fields in the human. So remember, we saw 4 tonotopically mapped fields in the cat. And here we have at least 4, perhaps 5 in humans. This one is labeled HG, that's Heschl's gyrus.

And that's probably primary auditory cortex. This is a view looking down on the superior surface of the temporal lobe. And this is what's called an inflated view. So it's like taking that cortex model and blowing it up like a balloon. And so the gyri are indicated by the lighter shading, and the sulci are indicated by the more dark shading. So that's what that is. These are the dimensions here. Posterior lateral is that direction.

So we have multiple tonotopically organized areas in human auditory cortex. Now, the paper that we read for today's class by Penagos et al talks about a center near A1, near Heschl's gyrus, which lights up in imaging studies when the subject is presented with sounds that have a strong cessation of pitch. So we talked about pitch a little bit earlier in the class. And this is the slide that I showed.

And it has a different title now. I think it was titled something like complicated sounds, or something like that. Well, a complex sound is simply a sound that has multiple frequencies. And so the stimulate used in the paper are complex sounds in that they have multiple frequencies. Earlier, we talk about this in terms of the context of musical sounds. Musical sounds almost always have a fundamental frequency, and then a whole bunch of harmonics. And to have a strong sensation of pitch, these musical sounds have a very tight relationship of the fundamental and the harmonics.

They can't be a random relationship. They actually have to be multiples of the fundamental. For example, this complex of tones, 100 hertz, 200 hertz, 300 hertz, 400, and so on, are multiples of one another. But if you had the fundamental be 100, the next harmonic be 150, the next harmonic be 190, the next harmonic be 230, where they're not multiples of one another, that stimulus would not have a strong pitch associated with it. These musical sounds are interesting because they have a strong pitch.

The pitch is almost always related to the lowest, or fundamental frequency of them. And the pitch is very invariant. As long as you have this nice pattern of harmonics that are related to each other by multiples, the pitch of this, no, which I think is a

piano note, and the pitch of this note, which-- let's do these two. This is a guitar sound, and this is an alto saxophone sound where the fundamental is the same. The harmonic amplitude is completely different.

But we recognize them as playing the same pitch. You can play a lot around with the amplitude of these harmonics. I drew them all the same. But clearly, they can be any jumble of pattern, as long as they're multiples of one another. You hear this as having the same pitch as that. And pitch is very invariant to things like where the sound is coming from. Pitch is invariant to how high and level the sound is. So pitch is a very fundamental attribute of the sound.

Defined in the psychophysics textbook is, pitch is that attribute of auditory sensation in terms of which sounds may be ordered on a musical scale. So this one's low, this one's middle, and this one's high. And so when cochlear implant users are programmed first, they take this electrode, and they stimulate it. And the user says, yeah, that sounds like a low one. Then they activate the next electrode. And the audiologist says, is this one higher, or is this one lower?

And if it's higher, then they route their speech processor, higher frequencies, into that electrode. So they do a pitch ranking in an auditory implants. Now, pitch of the complicated sound depends strongly on the fundamental frequency. Everybody knows that. Wow, you can play little tricks in these stimuli. You can do something like remove the fundamental frequency. How does that change the pitch?

Well, this guy becomes the new fundamental. That's what you might think. But actually, removing this fundamental is just like playing around with the amplitude of the higher frequencies. It doesn't change the pitch at all. And this is called the missing fundamental. And that's actually lucky for cheap speakers that might not have a very good base.

The fundamental is hardly there at all, but the piece still sounds musical, and it's not changed a lot. Or why does that happen? Well, this very nice multiples of 100 is still present. And so the temporal pattern of all these harmonics, if you add them up and look at this thing in the time domain-- remember, this was a graph frequency.

Whatever this looks like, it's going to repeat after 10 milliseconds.

Because it's period is still 100 hertz. I'm not a very good artist, but it's going to be the same. And it's going to be the same here. So each 10 milliseconds, it's going to repeat its pattern. It has the same regularity, even if you remove the fundamental. Now, you may not believe me. So let me give you a demonstration. And in this demonstration, there's a whole bunch of harmonics presented at first. And then in the second presentation, they remove the fundamental.

And you should listen to see if the pitch that your ears hear changes at all. On the second presentation, they remove this next harmonic, and so on, and so forth. And I think they end up with removing four different harmonics after they present the complete stimulus. So listen to this demonstration of the missing fundamental.

[AUDIO PLAYBACK]

-Pitch of the missing fundamental, or virtual pitch. You will hear a complex tone with 10 harmonics, first complete, and then with the lower harmonics successively removed. Does the pitch of the complex change? The demonstration is repeated once.

[SERIES OF PITCHES]

[END AUDIO PLAYBACK]

PROFESSOR: I think it's pretty clear. Does everybody want to discuss this? So when you lose these first 1 or 2, the pitch doesn't change a great deal. But by the end of the demo, this pitch is starting to sound a lot higher. So if you move some of these around or decrease their amplitude, this one's a low fundamental. The pitch doesn't change a great deal. Now, that is what they did in the paper that we read for today.

What they do is they have a complex tone with a whole bunch of harmonics. And they do a clever thing like they've done on this demonstration. They just select some of the harmonics to present to the observers in the imaging study. And the cleverness of this study is that by clever filtering of this harmonic pattern, they can

give you some stimulae that have really strong sensations of pitch.

Or in this one case, which I think is condition 2, a very weak sensation of pitch because of the particular harmonics they've chosen. So they have 3 stimulate. They giving you a strong sensation of pitch, and one that has a very weak sensation of pitch because of the clever way they've filtered it. And further clevering, all of these stimulae have the same regularity. They have the same regularity in terms of their temporal waveform.

So the cleverness of this study then is that the temporal waveform hasn't changed in terms of its regularity. But the subject's impression of the pitch, whether it's a strong pitch or a weak pitch, has changed. So by weak pitch, I mean something that sounds like a noise, or a click. Those stimulate don't have strong sensation of pitch because they're random. They don't have this nice pattern of harmonics. So I wasn't convinced by this verbiage and the figure.

So I decided I wanted to listen to these stimuli myself. And it was convenient, because I know all three authors. So Hector Penagos, when he wrote this paper, was a graduate student in the speech and hearing bioscience and technology program. Jennifer Melcher is a faculty member over at the Eaton-Peabody Lab. Her office is right next to mine. And Andrew Oxenham was a faculty member here at MIT, and has since moved to University of Minnesota.

So I started asking the authors, because they're all friends of mine, if I could have the demos. And one of them said, well, it's been a long time. I'm not sure I still have them. And the other author, is the second author I went to-- I won't say who it is, said, I got them right away. So he sent them-- that author sent them to me right away. And so I have them, and I'll play them for you.

Now, you'll listen to these stimuli. And what was surprising to me that I didn't get from the methods of the paper is that they don't just keep presenting the same thing over and over. The pitch actually moves around. And that's one of the nice parts of this demo is that you can actually tell that the pitch is moving around in the ones with strong sensation of pitch. Second thing that they did was, they added a little bit

of background noise to these. And it turns out that when you present a whole bunch of harmonics with the speaker, the speaker will introduce a little distortion.

Your ear introduces distortion. And they wanted to mask that distortion out. And the distortion is pretty low in level. And so this noise that's a continuous background is a pretty effective mask. I think you can still hear that these stimuli, in some cases, have pretty strong sensations of pitch. So I'm going to start out with condition one, which they say has a strong sensation of pitch.

And you can judge for yourself, whether you hear the pitch moving around.

[PITCH WITH STATIC]

PROFESSOR: Maybe it will go on forever. I don't know how long it would go on for. Well, anyway, could you hear those moving around? You could rank them. Here's number 2.

[PITCH WITH STATIC]

PROFESSOR: It's moving around, right? Number 3.

[LOWER PITCH WITH STATIC]

PROFESSOR: OK. Now, to me, those have strong sensations of pitch, I believe. I'm a believer. Now, here's the last one that they say has a weak [INAUDIBLE].

[PITCH WITH STATIC]

PROFESSOR: OK. So at first, I was expecting to hear no change in pitch at all. But actually, the pitches change a little bit. And so when you go back and say, it's a weak sensation of pitch, OK. Right? So we're believers? Or does anybody want to say-- all right. What happened in their imaging study? Well, this is pretty small figure. But it summarizes the results in that they had an activation of the circled area, which is near Heschl's gyrus-- just, I believe, anterior to it.

And this is the place that had a high activation. In the cases of the stimuli, was strong. Psychophysical sensations of pitch, but had a low activation in the case where there was a weak sensation of pitch. Other areas of the brain-- for example, Heschl's gyrus lit up for all of the conditions. And what's interesting about this study is they examined some of the subcortical nuclei that we've been talking about.

For example, the inferior colliculus and the cochlear nuclei. And these are the activations for those centers. Cochlear nucleus activates pretty much the same for all of the conditions. This black one is the one that was associated with the weak sensation of pitch. The inferior colliculus, perhaps a little bit less activation for that stimulus, but not significantly so. In the case of this center for pitched salience in the auditory cortex, clearly there's a lot less activation.

For that stimulus with the weak sensation of pitch, about the same amount of activation as you'd find with noise bursts there. So that, I think, clearly demonstrates, and other imaging studies have clearly demonstrated that this area is the region of cortex that becomes very active when we experience stimuli with strong sensations of pitch. In experimental animals, recordings from, for example, the marmoset auditory cortex.

You find neurons in an equivalent area that respond very nicely to harmonic complexes that have strong sensation of pitches. If you remove too many of the harmonics where the pitch changes a lot, the neurons fire less. And clearly, in those cases, remember, most neurons in the cortex are finally tuned to sound frequency. You can remove a whole bunch of these lower harmonics, and not change the response, suggesting that there is really signaling, that there is pitch associated with that stimulus rather than there's a certain kind of frequency.

So that's the bottom line for this study. Let me just mention a couple things that make auditory experiments difficult when you're trying to image the brain. Has anybody listened to an fMRI machine? An imager? Very loud, right? So you have problems with the subjects listening to the stimulus that you intend to present to them rather than listening to the imaging noise itself.

And so in this study, they went to great extent to try to reduce the imaging noise. The subjects were wearing protective earmuffs, and the stimulae were loudspeakers that led to those earmuffs in long tubes. Of course, you can't have a speaker right near the ear because there's a magnet associated with the speaker. So you have to have the speaker outside the imager. So they actually turned one of the imaging pumps off.

There is a lot of challenge in imaging such small structures as the cochlear nuclei in inferior colliculus. And they said to improve the detection of activation in these brainstem [? structures, ?] the data were activated using cardiac triggering. So does anybody know what that means? Well, when your heart beats, it pulses on all the arteries, and actually moves the brainstem.

Brainstem is so small, it can be moved. Cortex is moving too, but the areas are generally bigger. And so if you weren't going to take care of the heartbeat, the brainstem might be imaged when it was in this position at one point. And the next image might, when it was over here. So what cardiac triggering is, they record the EKG from the subject. And when they see the QRS complex, or whatever waveform from the EKG, and they say, that's the time to take the image.

So they only take the image at a certain point relative to the cardiac cycle. So the brain, even though it's moving, it's always moved in this certain position. So that's a challenge that's associated with imaging small structures like these brainstem nuclei. So then, another property, at least of this field of the auditory cortex, is to process stimulae that to have high pitch salience. So we've had two functions associated with auditory cortex then, just as a summary here.

One is processing stimuli with high pitch salience, and the other is processing sound stimuli that change in location. So those are the two things you can really hang your hat on, and what is done at the auditory cortex in terms of function. And for the pitch sensitive area, you have this area near A1. For the localization, you have A1, posterior field, and a field near the anterior ectosylvian sulcus, as we know currently.

All right. Questions? If not, have a good Thanksgiving. Don't eat too much, or enjoy eating too much, I guess. I'll see you on Monday.