

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.

PROFESSOR:

Last time we talked about processing of cues that are used for binaural localization of sound. Those being the interaural time and interaural level differences. And we have those cues because the ears are physically separated on the sides of our head. And because of the physical characteristics of sound, for example, the velocity of sound in air. And we talked toward the end of last time's lecture about the neural processing of interaural time differences in the Medial Superior Olive, or MSO. And we talked about the Jeffress model and how it can help recreate a neural mapping that has ITD sensitivity.

And that's the subject of the assignment for this year. So I just put the text for the assignment up here, just to mention that there was something added on at the end. But the first paragraph of the assignment is basically the lecture and the sketch of the Jeffress model. Except that these days, people are thinking that the axonal delay lines are not how the delay is created in the medial superior olive. And so I ask you for two other possible neural mechanisms that weren't originally postulated by Jeffress that could create delays. So that's sort of the heart of the assignment.

The second paragraph talks about some further updates to the Jeffress model. And there is a paper that is attached to last time's lecture. The paper by [? Brandt ?] et al where they discuss extensively some challenges, they call them, to the Jeffress model. And I might call them amendments to the Jeffress model, but things that weren't originally postulated by Jeffress that have come to light because of more recent experimental studies that don't fit with the original version and their updates.

And finally, we've been talking quite a lot about cochlear implants. And there's a very nice passage on cochlear implants in the textbook that I encourage you to read. And the final paragraph of the assignment is, what problems would cochlear

implant users, even those with a left and a right cochlear implant, have if they use the Jeffress model to localize sounds? So that's a little bit of thinking to solve that last problem, too.

But based on what we have talked about and what the textbook talks about for cochlear implants, you should be able to come up with an answer for that. So any questions on the assignment?

It relates very heavily to last time's lecture on the Jeffress model and the MSO. And I guess-- it doesn't say here, but three to five pages would be appropriate, I think. And it's due on December 4, which is the day of the lab tour.

So on that class day, which is a week from Wednesday, we'll meet at the Massachusetts Eye and Ear Infirmary instead of meeting here. So we'll send you a reminder. And I think now on the website there are directions to get to Mass Eye and Ear Infirmary. It's just a simple one-stop on the Red Line subway to get there.

So the assignment will be due then. And then I'll be able to grade it before we have a review session. And we'll talk about the correct answers for this assignment at the time of the review session, which is the next class after December 4.

So today, I want to march into new things. Now, we're going to talk about neural processing of interaural level differences. Remember if a sound is off to one side of my head, it's going to appear at a higher sound level in the ear facing that sound source than it will appear in the ear away from the sound source. So we talked about how big these cues were last time. And they can, for high frequencies, be maximally 20 dB in Interaural Level Difference, or ILD.

So those are processed, at least starting in the lateral superior olive, which is another nucleus in the superior olivary complex, which is close to the medial superior olive we've been talking about. So we'll talk about the process in the LSO. Then, we'll talk about projections of that nucleus and the MSO. And other parts of the superior olive to the next higher center in the brainstem, which in the auditory pathway is the Inferior Colliculus, or IC. So we'll be talking extensively about the

inferior colliculus, which is a large nucleus in the brainstem just caudal to the superior colliculus that you folks talked about extensively in the visual part of the course.

In the IC, you have some interesting neural responses that relate to something called time/intensity trading. I misspelled it. Not a very good speller, sorry. Which we'll talk about and define. We'll talk about some interesting characteristics of room acoustics, like reflections off the wall, and how those don't completely upset the apple cart of knowing where a sound is coming from because of things like the precedence effect.

And finally, we'll end up with auditory pathways in the barn owl. So this species of birds has very well-developed auditory systems. And some work at Caltech has shown that these animals have so-called "space" maps in a certain part of their brain. That is, there's a mapping of where the sound is in external space into space in part of their brain called the optic tectum. And we'll go over that.

And the reading for today is on how that space map is plastic and can be changed by experience. So we'll talk about neuronal plasticity of the space map. OK, so let's get started.

So the neural processing of the interaural level differences in the lateral superior olive. By contrast to what Jeffress cooked up for the MSO, this neural circuit is very simple. And here's how it runs.

The Lateral Superior Olive on the left side is here, the LSO. There's an LSO on the right side as well, but the circuit is shown for the one on the left side. This LSO on the left side gets input from the cochlear nucleus on the left side. That's excitatory input. And it gets input from the right side that's inhibitory.

And the way the inhibitory input works is the cochlear nucleus neurons on the right side project across the midline and into another sub-nucleus of the superior olivary complex designated here as the MNTB, which I think is spelled out here-- the Medial Nucleus of the Trapezoid Body.

So if you look at these sections under the microscope, there's a lot of crossing fibers here. And they sort of look, if you have a lot of imagination, like a trapezoid. And this nucleus is within those crossing fibers. So it's within the trapezoid body. That's how it gets its name.

And the MNTB neurons have inhibitory neurotransmitter. So the cochlear nucleus excites these MNTB neurons, but then they're inhibitory and they send their axons to the LSO. And they spill out or release their inhibitory neurotransmitter on to the LSO neurons. So how does this circuit work then if the sound is off to this side of the slide, off to the left side?

The sound will be of higher level in the left ear and it will exert a high excitatory effect on the cochlear nucleus here. And the LSO will be excited in a big way. Of course, a little bit of that sound is going to come over here to the right side, but it won't excite the right auditory nerve quite as much. And this pathway that eventually becomes inhibitory won't be as strongly activated. And so the inhibition will be less. So there's an interplay between excitation here and inhibition. And in this case with the sound off to the left side, the excitation will rule.

And so here's a plot of the firing of an individual LSO neuron that gets all these inputs. And if the sound is off to the left side, which is supposed to be this axis-- this is an axis of Interaural Level Difference, or ILD, where here on the left side of the graph the ipsilateral level is greater than the contralateral level. That case, the LSO neuron will be excited and it'll have a lot of response, a high amount of firing.

On the other hand, if the sound is over here on the right side, it will activate in a big way the right pathway, which will result in a big inhibitory input to the LSO neuron on that left side of the brain.

Sure, some of the sound is going to come over here and activate the left pathway. But that excitation won't be as strong. So in that case, the inhibition will rule. The LSO neuron gets a large amount of inhibition and its firing rate will be low. Its response will be low because it's inhibited. Now, what happens if the sound is right in the middle, 0 ILD?

Sound is coming straight ahead. The sound is the same at the two ears. And thus, the ILD is 0.

Well, it sort of depends on how you wire this up. Whether the balance is perfectly balanced. Then, maybe it would be a 50% response. In the case of this graph, it looks like the inhibition is a little bit stronger for equal sound on the two sides. And so the inhibition dominates at 0 ILD.

But in actuality, if you record from the LSO neurons, you find all sorts of combinations. Those that have 50% response rate at 0, those that have 90% response, and those that have 10% like this guy.

So this is a very simple circuit. No coincidence detection. No delay lines.

Well, you should kind of do a heads up here when I talk about delay and timing because this pathway coming from this contralateral side is a lot longer. The axons have to cross the midline. And then there's a delay here because at the synapse between one axon's terminal and the cell bodies of the MNTB neurons, there's a little bit of delay. This neurotransmitter has to be released. The MNTB neurons have to get excited and, finally, fire. So all that can take a half a millisecond or so.

It turns out this axon is a very thick one. This cochlear nucleus neuron here that provides this crossing axon is called the globular bushy cell. And we talked about that a little bit when we talked about the cochlear nucleus. It's not important exactly what type it is, but this has the thickest axon, really, of all the axons in the auditory pathway. So it gets across the midline very quickly.

And there is a synaptic delay here. So the contralateral input is going to come in a little bit later. So sometimes in recording some LSO neurons, you find a little bit of excitation from this side. And then right after, an instant later, a half a millisecond or a millisecond later, you find a big inhibition. And so that can happen.

The sum total though, in the case where the sound is off to the contralateral side, is a dominant inhibition. Now, we haven't talked about neural inhibition much in our

class, so maybe we should just mention it. And I should definitely mention the type of inhibitory transmitter that's used here.

So we have the MNTB neurons are coming like this and sending their axons to the LSO neurons. And they're inhibiting them. That's what I mean by this minus sign here. And so this is an inhibitory synapse which inhibits the LSO neurons. One can ask the question is, what is the inhibitory neurotransmitter?

And in this case, it's called glycine. So glycine is released from the MNTB neuron terminals onto the LSO neurons. And so how is that known?

Well, it's sort of a lot of little pieces in a puzzle here. The MNTB neurons themselves have the metabolic machinery that can make the glycine. They transport it down their nerve axons. Glycine is actually a fairly common chemical in all cells of the body, but these MNTB neurons are packed with it.

If you use antibodies to glycine, they stain these neurons much darker than most neurons in the brain. There are other glycinergic neurons, of course.

Their axons and their terminals are darkly stained. The lateral superior olive neurons have glycine receptors on them. When you put little puffs of glycine, which you can from a pipette. You can release glycine from a pipette in an artificial recording situation. When you puff glycine on to LSO neurons, they are inhibited greatly. So they certainly have the receptors for glycine.

There are uptake systems to take up the glycine after it's been released so that the inhibition doesn't stay on forever. When you stimulate, if you go in and stimulate electrically these MNTB neurons, you find the LSO neurons are inhibited. So a lot of little pieces of the puzzle go into the idea that these are glycinergic neurons. And this glycinergic input is very important in ILD sensitivity.

Now, a lot of textbooks will say the ITD sensitivity is created in the MSO and the ILD sensitivity is created here in the LSO. And we're done with it. But that's probably not true. This is such a simple circuit, you probably have other places in the auditory pathway where ILD sensitivity is also created.

For example, you have nuclei in the lateral lemniscus-- the pathway going up to the inferior colliculus. You have such circuits probably right in the inferior colliculus and maybe at other levels of the pathway. So this is not the only place where we find circuits for ILD sensitivity.

Now, we talked about last time how these ILD cues were prominent at high frequencies and very-- almost nonexistent at low frequencies. Because low-frequency sound can bend around the head very easily. So we had, I think, for 200 Hertz. Even for a sound source located directly off to the side absolutely 0 ILD.

For 6,000 Hertz, we have a huge ILD. So ILDs are not very important at low frequencies. If you go in to the LSO and record the frequency responsivity of the neurons there from measurements of their tuning curve. Now, remember what a tuning curve was?

Tuning curve, we had sort of over and over. It was a plot of sound frequency. On the y-axis is sound pressure level for a response. And we have the v-shaped functions and we picked off the most sensitive frequency. And that's the CF of the neuron.

Neurons in the superior olivary complex have beautiful CFs. And you could do a CF mapping, which is what's done in this study.

So in the LSO-- right here, you find CFs from 0 to 1 kilohertz. Right in this part from 4 to 10 kilohertz. In this part here, 20 kilohertz and up. And this is the typical kind of funny s-shape of the LSO that you see in coronal sections in this case of the cat superior olivary complex.

If you were to do this mapping in the cochlear nucleus, you'd find a lot of cochlear nucleus devoted to the low CFs, some to the middle, and some to high. In the LSO, you have a lot of the LSO devoted to the high CFs, which is where ILD cues are very prominent. So it makes sense that where you're processing ILDs, you devote a lot of neurons to responding to the frequencies where the cue is very salient.

The MNTB, which projects into the LSO, has a similar disproportionately large representation of high CFs. The MSO is just the opposite. There's hardly any MSO devoted to the very highest frequencies.

And remember, last time we talked about ITDs being ambiguous at high frequencies because the interaural time difference is still the same, but the sound can go through one or even more complete cycles by the time it gets to the other ear. And so you can't tell what ITD you're working with at these high frequencies.

On the other hand, where these ILD cues were weak at low frequencies, the ITDs are strong and salient. And there's a lot of MSO devoted to the low characteristic frequencies where the ITDs are prominent. So that's what this text means, that there's sort of a disproportionate amount of territory in these nuclei based on where the cue is important in the frequency domain.

Now, we have a little bit of an issue here with the LSO. This LSO on this side is going to respond to sound sources on this side of the body. And remember in most sensory systems, there's a crossing such that stimuli on the right side of the body evoke neural responses on the left side of the brain. This LSO has sort of got it backwards. It's responding with excitation to sound sources on the right.

Well, that's taken care of by virtue of the ascending projections from the LSO to the next higher center, which is the Inferior Colliculus, or IC. And that's diagrammed here in this next slide.

And this shows the projections of the LSO. Again, the LSO on the left side projecting across the midline to the inferior colliculus on the right side. And that projection would then predict that if a sound source was over on this right side and exciting the LSO neuron, that message would then get in a big way to the inferior colliculus on the left side. So that inferior colliculus would then respond to sounds on the right side of the body.

For a while, this field was a little mystified because there's also a projection from the left LSO to the left IC. It ended up being an inhibitory projection. So this projection

here that stays on the same side is mostly inhibitory. It's not exactly clear what that does, but it's there. You can sort of discount it in terms of the mapping of stimuli on one side of the body to responses on the other side of the body.

The MSO doesn't have such a problem, I'll just say in passing. The MSO, just because of its ITD map-- if you go back and review last week's lecture, you'll see that the MSO is already mapping ITDs for sound sources on the opposite side of the body.

Now, there's been a lot of work, especially in the early days of auditory neural science, on looking at the effect of lesions. And lesions are a little bit hard to do in a complex as hard as the superior olivary complex. Because if you go in and try to destroy the LSO, invariably right next door is the MSO on one side. And right next door is the MNTB. It's very hard to make selective lesions in the superior olive.

When you get to the inferior colliculus-- oh, there's a big nucleus. And you can go in and you can destroy it in an experimental animal on just one side. When that is done, and the animal is trained in a task to localize where the sound is coming from, it's very clear that a lesion of the inferior colliculus on the right side makes the animal unable to localize sounds on the opposite side of the body.

So lesion in the right inferior colliculus, the animal can't localize sounds on the opposite side. The animal still can localize sounds on the ipsilateral hemifield, in the side that you didn't lesion on because the other colliculus that's still there that's not lesioned can perform the task.

If you lesion both inferior colliculi, the animal can't localize sounds anywhere. But it's very clear that a lesion on one side of the auditory pathway here makes the animal unable to localize sounds in the opposite hemifield. So that's a theme for sensory processing that stimuli on one side of the body are mapped to neural nuclei on the opposite side of the brain. And that's very clear in the auditory pathway.

Now, I think last time when we had the demonstrations of ITDs and ILDs in headphones where we could present one and not the other, someone said, well, it

sounded a little bit like when we had an ITD when the sound was on the left side. And then later, we had an ILD with sound on the left side. They sounded a little bit the same, as if the sound could be put on one side by either ITD and ILD. And that was an interesting comment because of the phenomenon called time intensity trading, which you can do with headphones. I won't demonstrate it because it's so clear to explain it.

You can have someone listen in headphones and you can make-- let's see. An ITD such that when that's presented alone, the sound is perceived as if it's coming from the left side.

You can also present with those headphones the same sound, but make now in this case an ILD. So the sound is higher on the right side so that it sounds like it's coming from the right side. So now we have the time cues making you think it's coming from the left. The intensity cues making you think it's coming from the right.

When you put those two things together, which you can do artificially with headphones, you find in some cases that the sound sounds like it's coming from straight ahead. And this is called trading for time and intensity, or time/intensity trading. And you can balance one with the other.

If you do a lot of sound level on one side and just a little time on the other side, it sounds like it's off to the right. If you do them equal, it sounds like it's straight ahead. If you do a lot of time, a big ITD and a little ILD, it sounds like it's a little off to the left. It's a very clear psychophysical phenomena. Where do you find neural responses that correlate with time/ intensity trading?

Well, you find it in places like the inferior colliculus, where inputs from the LSO and the MSO first converge. So both LSO and MSO project up into the inferior colliculus. And here, you have the kind of responses that are shown on this graph.

This x-axis is now a time axis, except it's ITD. One side is the inhibitory lead. The sound is delayed so that the ITD is from the opposite ear. This is the side so that the time is delayed from the ear on the same side as the inferior colliculus you're

recording from. And clearly, this neuron is sensitive to that ITD if you vary it.

This is the percent neural response. 100% would be up here, 0 response would be up here.

Now, what's varied as the parameter here is the interaural level difference. So in this case, the contralateral ear is-- it looks like 35. And the ipsilateral ear is 45. They've kept the contralateral ear the same for the most part. And the ipsilateral ear level has changed. And clearly, changing the level also has a big effect on the neural response.

So here, for the first time in the auditory pathway, where we're finding ITD and ILD responses together in single neurons-- in a big way. You find them in a small way in the MSO and the LSO. But this is a huge effect here. And probably, there is the first place where you might have neural correlates of your perception for time/intensity trading. So that's one phenomenon I want to cover that probably has a neural correlate at the level of the inferior colliculus.

And here's another one. We've been dealing with very simple stimuli that have just, say, one ITD or one ILD. When you get into a room, it becomes much more complicated because of echoes off the walls and ceiling and floor of the room. And there's some very interesting experiments that are done with more natural stimuli like you find in rooms. And we'll go over just a few of them.

So here-- and this is an observer listening to a sound source off to the observer's left. And the direct sound is indicated by A here. And most of these data deal with the interaural time difference because this is a fairly low frequency.

This interaural time difference of the direct sound indicated by this big arrow in the air here favors the left ear. It's going to be arriving at the left ear first and the right ear a little bit later.

So if you just had that sound, the subject, obviously, would report that the sound is coming from the left side. And this is a plot of the sound for the left and right ear if you just had this arrow in A here.

And in green, it's the interaural time difference for just that direct sound. So you can see it's a fairly stable ITD and it's a negative ITD just by convention. Left ear is going to be negative here, ITD.

And this is an ITD in milliseconds of 0.4 milliseconds. And remember, we said if it was directly off to the left side, it would be about 0.6. If it were straight ahead, it would be 0. So what's the sound stimulus here?

Well, this is, again, changing things a little bit. We've been talking about very simple pure tone stimuli, or clicks. This paper is from a German group. And French or German speakers have this lovely speech sound, which is called like a trill or rolled R. And I, for the life of me, cannot do this stimulus. But I'll try. It's something like [ROLLING R]. As if you were to pronounce the word in German "reich."

Or in French, [INAUDIBLE]. It's impossible for me to do that because I don't speak those languages.

But anyway, here's the sound stimulus with all the echoes added in. And one of these traits just shows the left ear input and the right ear input. And these peaks here are the trills of the R. You can see them of just the left ear input as considered. So it's a trill of the R. And you can see how many milliseconds happened between each of those parts of the trill. Maybe like 40 milliseconds or so.

Now, when the subject is in a normal room, a lot more happens than just the direct sound. Some of the sound comes from the source, the R here, and bounces off the wall to the subject's right. Some of the sound goes beyond the subject's head and bounces off the wall to the right, and then the wall behind the subject and comes back there.

And this drawing here is when you take into account several of these reflections in addition to the direct sound. And the overall ITD is still plotted in green. And look what a mess it is.

The ITD is all over the place. It starts out like it should from just the direct sound. It's

negative here. But then real quickly, it goes past 0. And it goes way up here and it bounces around, then it goes back down again. It's all over the map. So where do you think that subject says the sound is coming from? Do you think the subject says, I can't tell, it's all over the place?

Well, you've been in rooms and you've listened to speakers in rooms. You're hearing me right now and there's a lot of sound coming off the walls to get into your ears on both sides. If you add up the reflections from the walls to the side and beyond, it's a lot greater than the direct sound in terms of the total energy. But you can close your eyes and you can know that I'm standing up here to this side of you. You don't have any trouble with that.

And if you just do a very careful experiment with this subject, the direct sound, and a couple of reflections like here. And you have the subject with a pointer. That subject will say the speaker is over there. It's on the left side. It's not behind or to the wrong side. So how do we do that?

Well, there is something called the precedence effect, which is very important here. Which helps you in a situation where you have lots of reflections. And what does precedence mean?

Well, precedence means something is dominating or something is the most important. And in these cases, if you look very carefully at the complex sound, when you add up all the reflections with the direct sound, you can see, if you have really good eyesight, that right at the beginning of the sort of burst of energy, the first thing to get to the subject ears is coming from the left side. So the left ear input here right at the beginning of these trills starts out. There's a bigger left ear input. And that's why you have a negative ITD right at the beginning of the trill before it starts going crazy.

It turns out that the most important part of this signal for localizing where the sound comes from is in the initial few milliseconds. So subjects bias their impression of where the sound is coming from right at the beginning of the sound. And they tend to suppress or ignore all the remaining parts of the sound, at least for localization

purposes. So what takes precedence here is the very initial part of the sound signal. And that's what the precedence effect is.

It's sometimes called by a different name when you're dealing with speech-- the Haas effect. But the precedence effect is more general for any type of sound. And it really means that when you're in a complex environment with all sorts of reflections, you pay attention to the very first, or initial, sound and ignore the others.

So the precedence effect can be studied carefully by just narrowing it down to a direct sound and one reflection by this kind of setup here. So here's a subject listening to two speakers in an anechoic room. And I think we've talked about anechoic rooms before.

Anechoic rooms have these baffles on the walls and the floor. This person is seated in a chair in a mesh, so he doesn't fall into the baffles below him, which are also absorbing sound on the floor of the room. There are baffles in the ceiling as well. So whatever sound this is presented goes right to the subject. If it goes beyond him, it goes to this wall over here and is completely absorbed. So there are no reflections. There's just the direct sound.

So you can say, well, one of these is the direct sound. And a little bit later, I'm going to introduce a second sound coming from somewhere else, which is the echo. So this might simulate a wall on this side of the subject that reflects. In this case, there's no reflection. But you can say this is an artificial echo presented by this second loudspeaker. What happens if we change the timing between those two sounds?

And that's what's plotted here. This is the delay between the first sound and the second sound in milliseconds. Now, for ITDS and sound localization, remember we've been talking about ITDs way down here from 0 to about 0.6 milliseconds. That's where you have sound localization way down here.

And when you do a delay that's that short, it sounds to a subject like there's just one sound. If the ITD is at exactly 0, this subject will perceive the sound source being directly in between these two speakers.

As the lagging sound gets greater and greater, the subject will start to perceive that the source goes to the original or first sound emitted. And when the delay comes to the maximal delay for the size of the human head, which is about 0.6 milliseconds right here, the subject will perceive that it's only coming from that initial loudspeaker.

Then, as delays get further and further on, up between about 0.6 milliseconds and a little over 10 milliseconds, the subject will still say, I still hear one sound. And it's still coming from that speaker. But now it's starting to sound a little bit different.

In fact, it doesn't sound as dead anymore. It sounds like I'm in a pretty live room. So what's a live room?

A room that's reverberant, like a church or a cathedral. It sounds roomy. It sounds like there's volume.

Over here, this is the region called the precedence effect, where you ignore that lagging sound. And it still sounds like there's just one sound, but it sounds different. This is the region of the precedence effect from 0.6 to maybe 10 milliseconds or so.

Then, as the delay becomes longer and longer, you start to perceive two sounds. You hear, let's say, a click from the first speaker and a second click a little bit later from the second speaker. And now, the delay is long enough so you actually hear an echo. You hear two sounds, an initial sound and an echo. So that's the perception. And there's this big region here called the precedence effect.

So I have a demonstration, if you don't believe me, about echoes. And the demonstration is really vivid, I think. The echoes become more-- you don't hear the echoes here because it's a precedence effect. Maybe it sounds a little bit more roomy, but you don't hear an echo. But there definitely are echoes there. You may not be able to perceive them.

So let me play this demonstration. The demonstration is-- I think, the best part of the demonstration is someone taking a brick and hitting it with a hammer. And that makes a big click, right?

Well, after that big sound, that impact sound, there's some echoes. You can't hear them very well except when they do this neat trick on this demonstration, which is to play the sound recording backwards. And then, the echoes start first, and then you hear the hammer hitting the brick.

They also have some text on here. And they read the text, and then they play the text backward. To me, that's not so obvious. So they do this demonstration of hitting the brick. They do it in an anechoic room right here first. There's no echoes.

Second, they do it in a normal room, like this room where there's some reverberation. But a lot of the reverberation is stopped by the carpet on the floor and the clothes I'm wearing. And the seat cushions absorb some of the echoes.

Then finally, they do this demonstration a third time. And they do that backwards, too. Then they do it a third time in a very reverberant room, like a church or a cathedral, where you hit the brick and it just sort of rings for quite a ways. So let's play this demonstration and see if it lives up to my description.

[AUDIO PLAYBACK]

-[INAUDIBLE] echoes. First, in an anechoic room. Then, in a conference room. Then finally, in a very reverberant space. You will hear a hammer striking a brick followed by an old Scottish [INAUDIBLE]. Playing these sounds backwards focuses our attention on the echoes that occur.

From ghoulies and ghosties, and long-leggedy beasties and things that go bump in the night, good lord deliver us.

PROFESSOR: OK, that's forward. This is backward.

-[SPEAKING BACKWARD]

PROFESSOR: OK, now the conference room.

-From ghoulies and ghosties, and long-leggedy beasties and things that go bump in the night, good lord deliver us. [SPEAKING BACKWARD]

PROFESSOR: All right, now in the reverberant room.

-From ghoulies and ghosties, and long-leggedy beasties and things that go bump in the night, good lord deliver us. [SPEAKING BACKWARD]

[END AUDIO PLAYBACK]

PROFESSOR: All right, so I like especially the sound of the hammer hitting the brick played backwards in the reverberant room because it's going to pssew. And all that pssew leading up to the impact is the echo that you just completely discount because of the precedence effect in the normal hearing.

OK, so that brings me up to the special part of my lecture, which is the reading. And this little quotation is by a musician. And of course, musicians love reverberant rooms, like churches or cathedrals, or concert halls, or whatever.

So E. Power Biggs, who was the organist at Harvard for a long time, made many famous recordings of organ music said, "An organist will take all the reverberation time he is given, and then ask for a bit more, for ample reverberation is part of organ music itself. Many of Bach's organ works are designed actually to exploit reverberation. Consider the pause that follows the ornamented proclamation that opens the famous 'Tocatta in D Minor.' Obviously, this is for the enjoyment of the notes as they remain suspended in the air."

So musicians love reverberations. And that's the reason that halls where we appreciate music, like Boston Symphony Hall-- and they have some measurements of reverberation time here for opera houses. 1.3 seconds. OK, that's the time the echoes take to decay.

Symphony Hall in Birmingham, 2.4 seconds. St. Paul's Cathedral in London, 13 seconds reverberation time.

Now in contrast, when you have theaters for speech, like Shakespeare drama theaters, you don't want all those reverberations. You want, for example, a theater for speech here is quoted at having a reverb time of 0.9 seconds because you don't

want all these echoes to interfere with your interpretation of the speech.

And average living room, 0.4 seconds. Reverberation time is given for the great outdoors. Anybody guess?

0.0. All right. The outdoors has no walls, ceiling, or floor. Now, why are we talking about this here?

Because in the inferior colliculus, you find some neurons that show precedence-like responses. And so here is a recording from an inferior colliculus neuron. And this was a study where they used infants and adults, but these are just the data for the adult. And the stimuli are two sounds. I think they are clicks.

The first sound starts at time 10 milliseconds. So this is the time axis. And this is the dot raster display. So each little dot here-- they're a little bit hard to see-- is a neural spike. And there are many, many trials. OK, perhaps 50 or 100 trials. And you can see reliably on all the trials, the neuron responded to the first stimulus.

When the second stimulus occurred at 101 millisecond delayed, the neuron also faithfully responded to the second stimulus. But as the delay was shortened less and less, the neuron eventually stopped responding for a delay that's-- in this case, it looks like about 20 or so milliseconds.

Well, this is certainly precedence-like behavior. That the neural response to the second stimulus is attenuated.

This delay, where this particular neuron starts cutting out, is not exactly where we stop losing the precedence effect in humans. This is from an animal. And this animal is anesthetized, so many of the processes are slowed down by anesthesia. And perhaps the animal is a little cool, which might make these things abnormally long. But certainly, this kind of precedence-like responses are on the way toward explaining the precedence effect at the level of the inferior colliculus.

And this is the work of Ruth Litovsky from University of Wisconsin. OK, so now I want to shift gears a little bit and go onto a different species, which is the barn owl.

And this is mostly the work of-- originally, Mark Konishi at Caltech. And now, Eric Knudsen at Stanford in California. And why did they choose to study the barn owl?

So many of us study mammalian models because we want to know what human hearing is all about. And it's hard to record from human brains, but we want to choose an animal that's like the human. So we use a mammal. Why did they choose the barn owl?

A lot of successes in neuroscience have been choosing an animal that's specialized for a certain task. And barn owls are very specialized for hearing.

So if you take a barn owl and blindfold the owl and turn a mouse loose in a room. As long as the mouse is making a little bit of sound, the barn owl can quickly fly over to the mouse and catch it and eat it. So the prey of barn owls are mice and insects, like grasshoppers that are down on the floor.

And if you've ever seen or watched a barn owl hunting, you can clearly see them at night. They don't come out during the day. But at night they come out. I watched one a lot of evenings when I was in California. The owl would come and sit on a basketball hoop. And it would just perch there. And it would move its head all around.

And I didn't really know what it was doing at the time, until later, when I read some of this work and it said that owls' eyes cannot move in their head. The owl's eyes are fixed. So this beautiful control of the eyeball position that we have in mammalian eye control is not present in most birds. And certainly is not present in the owl.

So to move eyes, you have to move the head. And of course, you're moving the ears as well.

In the barn owl, and in other words, you don't have an external pinna, which we said introduces a lot of help to localizing sounds. But the barn owl is a unique bird in that there is some external stuff going on, which is called the facial ruff. And that's formed by feathers on the face of the owl.

You can see they're sort of like fan-shaped down here. And over here, they go over the ear canal. There's an ear canal, certainly. They go over the ear canal in this perioral flap. And the opening for sound to get in is below that. And also, above it there's an opening. So the barn owl doesn't have a pinna, but it has some modified feathers on the front of its face.

Barn owl feathers are also interesting that they are modified. The feathers on the wings are modified. When this owl took off, when I watched it on that basketball hoop, every now and then it would take off and go down to the basketball court below. And you couldn't hear anything. Owl feathers are specifically designed so even the air going over them when the owl flaps its wings is completely silent because the owls don't want whatever it's hunting to hear them approaching it. So barn owl wing feathers are specifically designed to be acoustically silent.

OK, so a lot of work has been done on the barn owl pathway. Now, this is a little bit different because if you've looked at the brains of avian species, they're a little bit different. They evolve differently than mammals. But they have some analogous nuclei.

Here are the cochlear nuclei. That's with this little text here. This is supposed to be the owl pathway on the left side and the right side of the brain. This is the midline in dashed lines. The owl cochlear nuclei are split up into two parts. And the one that's featured here is labeled NM. So NM stands for Nucleus Magnocellularis.

And so we can all figure out what this means. Cellularlaris means the cell or the nerve cells. Magno means big. These are the big cells. So there's some other parts of the bird cochlear nuclei where the cells are smaller. But this is the big-- big cell part.

And in nucleus magnocellularis, you have beautiful phase locking. We talked about that being typical of the mammalian pathway. In the cochlear nucleus, the bushy cells have good phase locking. Maybe even better than the auditory nerve.

In the owl, the nucleus magnocellularis neurons have excellent phase locking. And

so they're keeping track of the stimulus waveform. The timing is important to them. They project centrally. The one on the left side and the one on the right side converge onto a nucleus that's sensitive to interaural time differences. That's the avian equivalent of the MSO. And it's called NL, Nucleus Laminaris.

OK, lamina means sheet. OK, and this is a sheet. It looks like that anyway in the anatomy. And there, the neurons are sensitive to ITDs. And there's a beautiful Jeffress model there.

Most of the papers on the mammalian MSO say, we know there's a beautiful Jeffress model in the avian nucleus laminaris. But in the mammal, we're starting to rethink it.

This is a beautiful Jeffress model where you find neural responses that are very strongly peaked to ITD. So they fire for a certain ITD. This is the firing rate. But don't fire much at all to other ITDs. They're strongly tuned to ITD.

The nucleus laminaris in turn projects across the midline here to the inferior colliculus. And we haven't talked about it, but there are several subregions of the inferior colliculus. The big one is called the ICC. And that's called the Inferior Colliculus Central part, or central nucleus. That's true in mammals as well. That's the big part. Some would call it the core. That's what's indicated here, the core of the inferior colliculus.

And it, in turn, projects to other places, like the lateral part. And it finally projects from the lateral part to the ICX. And the ICX is the Inferior Colliculus. And X stands for External. So the external part of the inferior colliculus.

And there is where some very interesting responses take place in the barn owl. And we'll look at those responses right now. So these experiments were first done by Mark Konishi in, I believe, the 1970s at Caltech. And the experimental subject is seen.

Here is the barn owl right here. And where are his wings?

Well, his wings are folded down. And he's in a little, sort of like a tube sock, if you will. His wings are constrained. And the tube sock is mounted on a pole or a pedestal. And he's just sitting there. He can move his head.

Actually, in many of these experiments, the head is clamped. But he's sitting there. He's awake. He's listening to the sounds.

And the sound is presented by a speaker. The speaker is on a big hoop. You see that hoop? And the speaker is sitting on the hoop. And there's a little motor in the speaker. And the speaker can be moved by driving the motor over here. Or you can move the motor the other direction. You can move the speaker over here or down here or up here, wherever you want to on that hoop the speaker can be driven by the motor.

And because the hoop is mounted on two posts on the side, the whole hoop can be swung up or it can be swung down. OK, so you can put that speaker anywhere this way and anywhere up or down that you want to. And so you can put that speaker in the entire-- any position you want to in the entire frontal hemi field of the owl.

And I suppose they didn't do that. I suppose you could put it in the rear hemi field as well. But these data are just from the frontal hemi field of the owl.

And since the owl's head is mounted and is not moving, you can apply a little local anesthetic and open up the skull. And you can advance an electrode into the ICX, the External Nucleus of the Inferior Colliculus, and make recordings from single neurons there.

And in this case, the recordings are made from the ICX on the right side. And what's found in this response plot here is these neurons have restricted receptive fields in space. So what's plotted here is the neuro-responsive field in dashed lines and the most vigorous part in the shaded area there.

And this axis, the x-axis, is azimuth. And the y-axis is elevation. And this is a fairly restrictive part of the whole hemi field. And one of these-- this is one response area for one neuron. And a whole bunch of these rectangles are plotted down here with a

little diagram of the owl right in the center of this globe.

So there's 1, 2, 3, 4, 5, 6, 7, 8, 9, 10-- about a dozen neural receptive fields for ICX neurons in that owl. Notice that they're all fairly discrete. That is, it only responds when the speaker is in a certain part of the frontal field of the owl.

Notice also that we're recording from the right side. And most of the receptive fields are off to the left of the owl. They're on the opposite side. So clearly, sound stimuli on one side are mapped to the opposite brain side. That's not true for these three, but they're close to the midline.

Most of the receptive fields are not straight ahead. They're actually down below the owl. There's one or two that are just a little bit above, but there are none that are way above. Most of them are down below the owl.

Remember, the owl is sitting on the basketball hoop and he's listening for targets down below. This makes a lot of sense that most of the receptive fields are down below the owl, not at the same level of the owl. It doesn't care about a mouse making a slam dunk at the other hoop. It cares about the mouse down on the basketball court.

And finally, maybe the most important and most interesting part of these responses is the progression of where the receptive fields are in space versus where they are along the dimensions of the ICX. So that this receptive field was located over here in the ICX. And as you move this way and encountered a different neuron, it was located over here and its receptive field was moved over this way a little. As you move further in that direction, you encountered another receptive field even further in this dimension. And then finally, way over here laterally, you encountered these receptive fields that were way off to the side.

So along this dimension of auditory space, receptive fields were found along this dimension of the ICX. There was also a mapping going this way. This is clearly then what some people call a space map. A mapping of auditory space into position within a certain part, the ICX, of the owl's brain.

OK, it's a beautiful mapping. Neuroscientists love mappings. You've heard me talk about tonotopic mappings out the wazoo. We love tonotopic mappings because CF is then very important. This clearly is important to have organization of the receptive fields on a dimension in the brain.

People have spent many, many years looking for these kind of mappings in the mammalian auditory pathway and not found them. People have looked in the inferior colliculus, in the analogous part, the central part, the external part. It is hard to record from the external part of the mammalian inferior colliculus because it's a small edge structure, but it has not been found.

People have looked in the medial geniculate, in the auditory cortex, and looked for organization of spatial receptive fields and not found them. So on nuclei in the main part of the auditory pathway, you do not find space maps in the mammalian system.

So the one place you find spatial organization in the mammal is in the mammalian superior colliculus. And you're probably going, huh? I thought you just said the superior colliculus is visual?

Well, it is. But it's a layered structure that, if I'm not mistaken, the top three layers are exclusively visual. But if you go down to lower layers, the bottom layers of the superior colliculus, you start to encounter neurons that respond to visual as well as auditory stimuli.

And you may have talked about the visual mapping of the superior colliculus. And those neurons in the deep layers that are also responsive to auditory stimuli-- they're mostly ILD sensitive. They're mapped in line with the visual receptive fields. They're also space mapped to a certain extent.

Now, that nucleus clearly is not on the main drag of the auditory pathway. The auditory pathway is cochlear nucleus, superior olive, inferior colliculus, medial geniculate and cortex. So you do have a space map in the mammalian deep layers of the superior colliculus, but not on the main parts of the mammalian auditory pathway. So that's been the finding.

Now, that's negative evidence. It's not clear that I won't be back here teaching the course next year. And we'll read a paper that says, ah, space map in wherever. But it hasn't been found so far with this one exception.

Now, the paper that we read for today's class talks about-- back to the barn owl. A place in the barn owl, which is-- they call it the optic tectum in the birds. But it's analogous to the superior colliculus in mammals. A place where, as I said, you find auditory spacing maps that are in line with visual space maps.

And they do a very interesting and elegant experiment where one of those maps is distorted. And you study the resulting effect on the other map. OK, so how did they do that?

So this is an owl, but it's a juvenile owl. An owl chick. And it's wearing some interesting things on its eyes. Those aren't its eyes. Those are some prisms that the investigators have put on the owl's eyes. And they deflect the visual field a certain amount depending on the size of the prism.

And I can't remember what the deflection was. I seem to remember 30 degrees. So the visual field is deflected 30 degrees. And as I said before, the owl's eyes are fixed in the head. So putting on these goggles, no matter what-- the goggles are going to move if the owl moves its head. So no matter what the owl does, the visual receptive fields of all these neurons, everything in vision is shifted 30 degrees.

This is normal. These might be receptive fields from neurons in the brain somewhere. This is when the prisms are added.

Here, you've shifted the visual receptive field. The auditory receptive field-- you haven't changed the ears at all. The auditory receptive field is the same.

What's found when you do that in juvenile owls, you come back eight weeks later. You find, oh my gosh, the auditory receptive field has actually moved. It shifted. You knew where you were recording from in the brain. In this case, the recordings are made in the optic tectum, in the superior colliculus.

You know the dimensions and you come back and you're expecting to see auditory receptive fields like that, but they've been shifted. So juvenile owls with prism experience, given a number of weeks to compensate, shift the auditory receptive fields so that the two are back in alignment.

This group has also shown that if you do this experiment with adult owls, you don't get such shifts. You come back to an adult owl with these same prisms, eight weeks later you have still a mismatch. So plasticity clearly takes place in the juveniles and not in the adults.

And they've then likened it to those old folks of us trying to learn a foreign language. It's really tough because we didn't have experience with it while we were juveniles. So neural plasticity and these learning new things-- if this is learning something, it's a bit of a stretch. But it's more difficult to learn things as an adult than it is juveniles.

Now, the experiment-- an even further twist in this paper that we read for today, which is now the subject is adult owls. The recordings were made from adult owls. But there are two groups.

One is a plain, old control adult owl. The other is an adult owl that when it was a juvenile had experience with the prisms. That experience was long ago, six months before. The prisms were on for eight weeks or so. These shifts took place. The prisms were removed. The owl is allowed to grow up and become an adult with normal vision.

Then, take these two groups of owls, put the prisms on again. We've already said that adults don't have the capacity to remap. So many of them just stayed the same. These are the open circles here. The auditory receptive field is the same. This is before the prisms. This is eight weeks after. The open circles are from the adult owls that didn't have any juvenile prism experience.

But the adult owls-- they're now adults with the prisms on. The adult owls with the juvenile prism experiences, one of the neurons is recorded here. It has an auditory receptive field that's shifted to make it in line with the shifted visual receptive fields.

This is now showing then that adults that have the juvenile experience have some plastic ability to re-map their auditory receptive field so that you have alignment with the visual receptive fields. Clearly, a very interesting experiment. Showing, if you will, a type of maybe memory trace that these owls have retained something that was altered by the juvenile experience. So what could that be?

The group has in other studies gone on and looked at the projections between these two boxes here from the central nucleus of the colliculus to the external nucleus. And clearly, shown in juvenile animals that have the prism experiences that the axons that were headed for direct innervation take a little change and they regrow some axonal projections. They've studied these.

And it looks like this regrowth of axon between those two areas. And that's what was meant by this little circle that says locus or site of plasticity is manifested by a change in growth of axons right there.

And maybe those adult animals that have juvenile experience retain some of those axonal projections that have been changed as a result of the experience. And clearly, it takes many weeks for axons to grow or change their connections. Maybe that's the thing that's much easier to do if you're a juvenile animal and you're reacting to these changed stimuli.

Another way to change responses there is to have everything connected to everything, but certain things emphasize certain synapses very strong and the weak ones not emphasize. Maybe when you have the prism experience, the previously de-emphasized synapses become upregulated without any change of axons. But clearly, this group has shown that the axons have changed their growth patterns. So that could be a mechanism for the plasticity.

And I think that's all I wanted to say today. So I have five minutes if you guys want to ask questions about anything. Yeah.

AUDIENCE: Can you go over how prisms actually change auditory perception again? Like, how are the prisms--

PROFESSOR: Back to precedence? Is that what you're--

AUDIENCE: How are prisms changing--

PROFESSOR: How are prisms?

AUDIENCE: Yeah.

PROFESSOR: That's not clear. What's clear is that the fields are mismatched.

AUDIENCE: So we don't know why that happens?

PROFESSOR: That's correct. Yeah. But I presume what is happening to the owl during these eight weeks is the owl is seeing an object and hearing it as well. For example, a mouse down here. The owl sees it and it goes for it. But actually, because its visual fields are off, it goes over here and the object is over here. But the auditory cues, if it paid attention to them, it would go here. So it's sensing a misalignment in experience.

There's no training involved here, but the owls are allowed to hunt prey and experience environmental sounds. So they clearly then have a mismatch between vision and audition in these eight weeks. Yeah.

AUDIENCE: Does this sort of then, I guess suggest that the auditory input is somehow more important? Because rather than say a visual input shifting [INAUDIBLE]?

PROFESSOR: Yes. Yes, you could say that. I mean, it would be interesting to do the converse experiment. Keep the eyeballs normal and somehow distort the auditory receptive fields. So you could do that with ITDs by putting a tube and lengthening the ITD on one side. That would be an interesting sort of counter experiment to get at what you're asking about. It would distort the other cue.

OK, great. We'll see you back on Wednesday then.