HHF Research Webinar Transcript
Selective Attention Across the Auditory System | January 22, 2024, 5pm ET
Ross Maddox, Ph.D.

ANIL LALWANI - I've realized that I'm never going to host Saturday Night Live. So I thought I'd begin by saying, live from New York. Welcome to the Hearing Health Foundation Research Webinar. Welcome. I'm Anil Lalwani and I appreciate your joining us today. And this event has a live captioner and is being recorded. You can enable closed captions by clicking the CC button in the bottom of your screen, the toolbar at the bottom. And if you need any other assistance using Zoom, please follow the link to the technical guide shared in the chat.

Now humans are remarkably a depth at listening to one person while in the presence of substantial background noise akin to being at a cocktail party, restaurant, or even at home. Now what does this look like in the brain? In this research webinar, Professor Ross Maddox, will discuss how we can measure the electrical activity while individuals are listening to a natural speech to kind of get a glimpse of how selective attention impacts neural coding of speech, both in the cortex, but also lower levels of the auditory system.

By way of introduction, my name is Dr. Anil Lalwani. I'm a professor and Vice Chairman for Research in the Department of Otolaryngology Head and Neck Surgery, as well as Associate Dean for Student Research at Columbia University, Vagelos College of Physicians and Surgeons in New York. I'm also a board member of Hearing Health Foundation and the Head of HHF's Council of Scientific Trustees, which oversees the Emerging Research Grants Program, also affectionately known as ERG. Now ERG is a competitive program that awards funds to researchers conducting cutting edge hearing and balance research, much like you're going to hear today. Now these grants supported many leaders in our field to become successful scientists, including our illustrious speaker today.

Dr. Maddox was also an ERG scientist in 2013 and was generally funded by Royal Arch Research Assistants. Now this month, Dr. Maddox became an Associate Professor at the Kresge Hearing Research Institute at the University of Michigan, my alma mater and the National Champs in football. So go Blue! Now prior to that, from 2016 to 2023, he was on faculty in the departments of biomedical engineering and neuroscience at the University of Rochester.

The ERG program that provides seed money to scientists just starting out in their field of research, is not only possible through the generosity of supporters like you, oh, it is possible to the generosity of supporters like you. If you'd like to support our work on hearing loss, tinnitus, and related conditions, you can do so today by going to hhf.org. That's Hearing Health Foundation, hhf.org/donate.

Now without further ado, we'll move to Dr. Maddox's presentation and please do ask the questions through the Q and A box linked at the bottom of the screen that we'll try to answer following the presentation. Dr. Maddox and a fellow Wolvorine, take it away.

ROSS MADDOX - Thank you. All right. Thank you for that nice introduction. I'm very happy to be here today to talk about some of the work, I and my lab have been doing recently. So the talk is going to be called Selective Attention Across the Auditory System and the "across" there is the really important part. We're going to try and sort of look, take a really broad view of where these processes are happening. But I want to start by specifically thanking Hearing Health Foundation.

As Anil mentioned, I had a grant from HHF in 2013 that grant was awarded after being twice rejected but for a postdoctoral fellowship by the NIH. And it really, it's hard to overstate the difference it made getting that. Since then I have been fortunate to review a lot of Emerging Research Grants and it's

always a pleasure to see some of the amazing work that's coming out of often sort of newer faculty and postdocs. So it's really a great organization that's supporting great work.

So the outline for today, it's going to be in four parts. We'll talk about what selective attention is, particularly in noisy places. Then we'll talk a little bit about how we study, what's going on in human brains. Then we'll talk about how we do that specifically with continuous speech like you're hearing right now. And then we'll get to this across the auditory system idea for how we look all over the place. Look at these responses all over the auditory system.

Okay. So part one, imagine, if you are at a noisy bar, a noisy restaurant, or maybe you don't have to imagine, maybe you're there right now, but this is a very challenging listening environment. And what you're faced with is a problem and something that is often called the cocktail party problem. And basically the problem is this, you receive a mixture, sorry to take the analogy a little too far, but a cocktail is a mixture. And so what you're receiving is an acoustic mixture and your brain's job is to separate the mixture into its ingredients and then choose which one it wants to pay attention to.

I'll do a little sound demonstration now. Let me just make sure my sound, it looks like it is sharing. Good. Okay. So what you're going to do is you're going to hear two people talking. One is me. And I'll be saying something that starts like this. The best part about a cocktail party. So I'm going to be saying what the best part about a cocktail party is, and you need to listen and ignore the other talker and listen to what the best part is.

Here's the mixture: The best part about a cocktail party [Wife] It's not the drinks, it's the company - is the cocktails. So you probably heard that I said the best part is the cocktails. Now you may not have been able to do that and that sort of underscores the importance and also the delicacy of these processes.

But if you were successfully able to do that, I would ask you a follow-up question which is, what did my wife say the best part of a cocktail party was? I'm unfortunately not in a room where I can ask people to raise their hands, but pretty much always no one is able to say so.

I'll play you the mixture again, but now listen to my wife what she's saying. The best part about a cocktail party - [Wife] It's not the drinks, it's the company - It's the cocktails, It's not the drinks, it's the company. And of course, she's right. So our brains are so adept at paying attention to one sound over the other that you often don't. You don't even get anything from the message that you were not paying attention to.

We'll leave our cocktail analogy behind for a moment and I'm going to show you what this actually looks like when we look at the sounds. So what you see at the top here are two what are called spectrograms, and these are visual ways of representing sounds. And you have time going from left to right on this horizontal axis and frequency on the vertical axis. So it turns out when we speak, we're making all sorts of frequencies at all different times. And so this one on the left, is me saying the best part about a cocktail party, et cetera. And the one on the right is the other talker.

But we don't receive these messages separately when they're spoken at the same time, instead we get this mixture. And you can see even visually, it seems like it should be easy to tell which of these lines is coming from which message, but it's actually not. It's quite difficult. Even so our brains are able to separate that mixture into its two separate parts and then choose which one to listen to.

This phenomenon is not newly studied. There's a very seminal paper from 1953 by Colin Cherry. And in that paper, there's one particularly important passage where he coins the term cocktail party problem. And this problem is defined as how do we recognize what one person is saying when others are speaking at the same time. He then goes on to give several ways that our brains might accomplish this, which I'll sort of summarize like this.

Essentially when you are listening to speech, especially in noisy locations, you're not just listening to the sounds. Instead your brain is basically estimating what the sound you want to be listening to is, and it's making that estimation in a bunch of ways. It's using the location. If the two people who are talking or lots of people who are talking are coming from different places, then your brain may hone in on one particular location.

If the person you're listening to has a low voice or a high voice relative to the other people, then your brain will use that pitch. All people have different timbers which is just sort of the way a sound sounds for lack of a better definition. We also often use visual cues to understand speech, which is a very clearly not an acoustic phenomenon, but it is a speech phenomenon. And then there's things sort of more related to language like the context, what is the topic at hand and expectation.

The best example of that is that I can stop talking mid-sentence, but you still know what I was going to say. So selective attention, it is an essential function or maybe the essential function of the auditory system. But Cherry's study in 1953 wasn't a study of the brain, at least not in measuring brain signals. It was a behavioral study.

So how does this happen in the brain and what areas of the brain are involved? And this is worth studying for a couple reasons. One, it's this really important process and we need to understand at the basic level, but two, this is something that a lot of people struggle with for a variety of reasons and we need to understand this, how this process works successfully to also understand disordered listening.

Okay. So how can we actually measure responses to sounds from the human brain? So here is a human brain and I like to think of this sort of the pink part on top. This is called the cortex, kind of a wrinkly boxing glove is how I typically describe it. And the auditory cortex is kind of on the thumb of that boxing glove. And we can put electrodes on top of the head and these are called EEG electrodes or electroencephalography. And we can actually measure voltages that come from all over the cortex, not just the auditory cortex, but the cortex isn't the only part of the brain that's involved in processing sounds, in fact not by a long shot.

It turns out sort of on the inside of the brain, we have what's called the auditory brainstem. And I'm going to use the word auditory brainstem a lot or the phrase auditory brainstem. I'm also going to use the word subcortical a lot. Subcortical is obviously in contrast to cortical, and sub means below. So this is what's happening to the sounds before it gets to the cortex. So what is electroencephalography?

Basically when we say brain activity, this involves charged particles called ions, essentially moving around and when charged particles move around, this is an electric current. This is an electric phenomenon. Although it's very, very small. If you get enough neurons, which are the cells inside the brain active at the same time. So they're moving ions back and forth. If you have enough of them and they're lined up spatially, then these create a measurable voltage outside the head. And this was first recorded in humans by Hans Berger in 1924. And as I was looking at this slide, I realized this is exactly 100 years ago. So we can't sing because this is a webinar, but I would like to wish EEG a happy 100th birthday.

There's a specific kind of EEG called evoked potentials. Potential is another word for voltage. And the way an evoked potential works is you present a stimulus, evoking a response, and then you can measure that response. And so you might present a stimulus and get something that looks like this. So sort of voltage that kind of goes up and down and then settles down for a bit. But the brain is really noisy. And by noisy here, I don't mean acoustically noisy, but they're actually, the cocktail party is a reasonable analogy.

If the part of the brain you care about is doing this sort of up and down bit, the rest of the brain is also recorded by these electrodes and it's doing all kinds of stuff that might be, or probably is, totally unrelated to the stimulus you played to the sound you played.

What we do to record evoke potentials is we present the sound over and over again and record the voltage that follows each time. And then what we can do with all those responses, all those individual trials is we can average them. And when we do that, we get what looks, what we can see here on the bottom, this average response. So while the actual response which is indeed buried in all these individual trials is very difficult to see, when you do it multiple times and average those, then you can start to see that the noise level goes down quite a bit.

Now this average response, it's not exactly the stimulus response you see here, and that's because there's always noise. But we can reduce the noise a lot by repeating a stimulus multiple times. And the more trials you collect, the less noise you'll have. So how can we use evoked potentials to study

attention? There is a very, very nice paper from 1973 by Stephen Hilliard and others. And what they did is basically, they presented the exact same set of sounds multiple times, but they changed the instructions to the listeners. Sometimes they said, okay, you're going to hear stimuli in both ears. Pay attention to the stimuli coming on the right. Sometimes they said, you're going to hear the same stimuli again, but now pay attention to the ones on the left.

So the only thing that differed between conditions was which one is the listener paying attention to? And what they found is that the right ear cortical response was bigger when attending to the right and the left ear cortical response was bigger when attending to the left. And you can see that in that figure here. So here's the right ear stimulus and the solid line is a 10, right. And you can see that this has a bigger excursion from zero here than the attend left, which is the dotted, but that's reversed when you pay attention to the left stimulus.

Now this was done, we mentioned averaging. So they had to present the same stimulus and what they used were something called tone pips, which are sort of just tiny little beeps. They did this with a device called a signal averager, which at the time cost $9,000 and these days would cost $86,000. But that averaging, this was a relatively new technology at the time. And that averaging was the only way really that you could get any sort of readable evoked potential. But tone pips are not sort of a very natural stimulus.

And here I actually want to go back to Cherry's paper to the abstract. And he points out that in this study that rather than use steady tones or clicks, these artificial stimuli, they used continuous speech. And that's really important because continuous speech, that's what we listen to, that's sort of the most important stimulus for human listeners.

So basically, what sort of a better experiment that's sort of, or a more natural experiment would be one that replaced artificial stimuli like clicks, which sound like this with speech, which sounds like this: [Narrator 1] Chapter one, a floating reef.

So now we get to the next part of the talk, which is how can we measure brain responses to continuous natural speech? And what that'll allow us to do is create really natural experimental paradigms where people are engaged in a task that is much more commonly, what's just something they're used to, paying attention to someone talking basically.

And so, a sort of typical evoked potential experiment would look like this. We have a person, that person has a brain and we are measuring the signals from that brain. And again, as a reminder, I won't force you to listen to them again, but you could play a whole bunch of clicks and you could measure the response to those clicks and average them. And you would get something like this, some evoked potential like this. I'm an engineer by training and engineers like to approximate things as black boxes.

So I'll approximate that as actually my head and my brain. I will approximate them as a black box. And so basically, you can approximate that auditory system as a black box that takes in clicks and turns them into evoked potentials that look like that. But again, that's using clicks. So is there a way where we can basically take speech, a natural sounding speech, and when we do that and measure the brain response to it, because it isn't clicks, we won't get nice, well-defined evoked potentials. But is there a way we can take speech input, measure this really messy output and then figure out what the evoked potential would be?

That's sort of the idea here because this evoked potential is nice, it lends itself well to interpretation. Just like the evoked potentials from the 1973 Hilliard paper did much more so than these sort of kind of difficult to interpret ups and downs. And so this is something when you do this, when you present a natural stimulus, record the long EEG in response to it and then mathematically figure out what the evoked potential would be. This is called a temporal response function or a TRF. And basically presenting these stimuli and figuring out what the click evoked potential would be.

A good analogy for this is, if you've ever gone into a concert hall or a cave or any room that has sort of a long reverberation. I always sort of get the urge to sort of clap my hands and listen to what the room sounds like. When you do that, you're measuring something acoustically that's called an impulse

response, because your hands when you clap create an impulse. They sort of very short sound and you can sort of get a sense for what that room does.

But you could also walk into a room or a cave or whatever and speak. And if you spoke and listened to how that room sounded, you would actually be able to guess pretty well what it would sound like if you had clapped your hands. And that's basically what we're doing here, but rather than a cave, we are doing that with the auditory system. So essentially, and I apologize, this slide is a little technical but don't worry too much about it if it's not working for you.

So we take a stimulus, we record the EEG data out, and we want to get this response, this TRF. We usually though we have to transform the stimulus and that's something called a regressor. I'll talk a bit more about that in the next slide. But so the stimulus is something that you present. The EEG data is something you record. The regressor is something you as the experimenter choose and the response is something you calculate. And when I say you choose the regressor, what I mean basically is you take the stimulus, the sound you played and you do something to it. Again, I promise, I will very soon explain what I mean by that.

So the TRF that's computed through something that's called deconvolution, this is just a mathematical process that allows us to do that. And the regressor as I mentioned, is the sort of transformation of a stimulus and it actually matters a lot to what types of responses you get. You could take the same stimulus and the same data even, and if you change the regressor you'll get a different response. So here is a short segment of speech, I can't remember exactly what it is. I think it might be a, no, it doesn't look like an utterance of Thursday. Anyway, I can't remember what's being said.

But a typical evoked potential paradigm would just sort of treat the onset of this speech as time zero and then say, well, the evoked potential sort of did this once the speech started. But the TRF works differently. It takes this regressor. And so here's where I can show you that. For the original TRF experiments what they did is they computed what's called the envelope of the speech and the envelope is this red line, which is basically how big is the speech stimulus right now.

So this is sort of the slow variations in amplitude where speech also has these really sort of fast fluctuations. And these fast fluctuations are what determine things like the pitch of the person talking. What the TRF does then, is instead of treating just the onset of this speech and measuring the response, it sort of says, well, this envelope when it's bigger, will probably be evoking a bigger response and when it's smaller, it'll be evoking a smaller response, but it is kind of at all times evoking these responses. And then what we can do is try to untangle that. Excuse me.

And this is where deconvolution comes in. So the regressor I was talking about for cortical responses is that slow envelope, but that's just for one person talking, which isn't really the best way to study attention. So what you can also do is actually present two people talking. So you can have one speech stimulus and another speech stimulus. And you can then compute the envelope for the first stimulus, the envelope for the second stimulus. You can record the EEG data, which of course, will reflect the responses to both.

But because you know what the stimuli were, because you're the experimenter, you can separately compute the responses to each stimulus. And if you're doing that in an attention experiment, then you can tell people just like Hillard did with the tone pips, pay attention to this person talking or pay attention to this other person talking. Or if you're using audio books, which is commonly done, pay attention to this story and ignore this one or vice versa. So how does this look in the cortex when we use continuous speech?

So Power et al. in 2010 measured the TRF to two simultaneously presented speech streams just like I just showed you. And they found really nicely similar results. Basically what they saw is, so these are now not technically evoked potentials in the sense that they're not sort of onsets. These are the TRFs for the continuous speech. And they found here the solid lines are the attended stimuli and the dashed lines are the unattended. And you can see in both cases for both stimuli, the solid lines are bigger than the dash lines. Here what we have are called scalp topographies.

And this basically just shows that the attended speech went more, the attended TRF was more negative kind of on this part of the head, which is where auditory cortex tends to show up even though

it actually exists more on the sides. And you can compute the difference between the attended and the unattended, and there is a rather large one. So attention's impact on these cortical responses. This was an early study that showed it for continuous speech.

It has been studied quite a bit more and people have done really clever things, looking not just at the envelope but looking at various linguistic features, looking at audio visual integration, using other types of ways of measuring the brain such as electrocorticography, which is where you actually have electrodes on the cortex. All sorts of work has been done in the last, I don't know, I guess 14 years at this point. Okay. But so far, everything I've told you about has been cortical. Okay.

So selective attention we know allows us to focus on this one signal while ignoring the unwanted ones. And the auditory cortex shows larger responses to attended speech then ignored speech. But a lot of neural processing happens in areas before the cortex in these subcortical areas. So I'll get more specific in a second, but just to remind you, here's the cortex which EEG mostly records. Here's the auditory brainstem, which is again kind of tucked up into the middle of the brain. It's a bit far away from those electrodes, but we can actually measure responses from these subcortical areas and that gives us something kind of uncreatively called the auditory brainstem response or ABR.

What the ABR is, it's is an evoked potential again, and it shows early brain activity in response to really short sounds, so like clicks for instance. And talk about averaging, it takes a lot of averages. It's typically several thousand averages to get a usable ABR. The ABR's main use is actually hearing screenings and hearing diagnosis when behavior isn't possible. The biggest group of patients for this are infants. And what's really interesting about it, is the component waves, and I'll show you some component waves in a couple slides. They're generated by very specific parts of the subcortical auditory system.

We've talked a bit about cortex and like I said, that's not all there is. So sound is first encoded in the sensory periphery. So the cochlea is what turns sound from physical vibrations into neural electrical neural impulses. Those then travel by the auditory nerve to the cochlear nuclei. This goes to a few areas, one called the superior olivary complex, one called the inferior colliculus. And I'll just point out here that this is a very simplified version of the auditory system. It's quite a bit more complex than this.

This then goes to the auditory thalamus and then it eventually makes it to the auditory cortex. So there is a lot that happens on the way up, but the way up isn't even the whole story. It turns out a full half of the neural connections in this auditory pathway are actually downward. And by downward, I mean they go from the auditory cortex where sort of more higher order cognitive processes are happening back down to these subcortical areas.

The auditory cortex projects to many of those subcortical areas. And higher subcortical areas like the inferior colliculus project to lower ones. And you even get projections down to the cochlea itself. So there is a lot going on here. Sound is passed upward and processed and information is extracted at these sort of ascending levels of the pathway. But there's also these downward connections. Another word for that is efferent. There are these efferent connections that are a bit more mysterious.

What is the ABR? The ABR, it's an evoked potential. It looks like this. Now to the uninitiated, this probably just looks like a bunch of squiggles, but this is actually across people. It's fairly stereotyped. And what's really cool about it is that it shows very specific waves. So we have wave one, wave three and wave five are sort of the most repeatable and these waves correspond to specific parts of the auditory system. So wave one corresponds to the auditory nerve. Wave three comes from the cochlear nucleus and wave five comes from a mixture of sort of higher subcortical areas.

The ABR, it's great, but it's got some challenges to recording, especially in response to speech. The first is that it is very fast, the second is that it is very small. Why is very fast an issue? So here is a cortical TRF and putting it on the same timescale, here is an ABR. All those waves I just showed you basically happen before the cortical TRF even starts. And if this is our speech waveform, our speech stimulus with our envelope on it, that envelope that we use to generate the cortical TRF, which fluctuates at about the same rate, is no use for looking at the ABR.

So we have to do something else. And what is not really obvious actually. So with cortex, we know that we put in a speech stimulus, we do the envelope. But the envelope is very slow, the subcortical

response is very fast. So how can we measure the ABR to continuous speech? And what we ended up doing is we took the speech and we re-synthesized it.

And what I'm going to do first is just play the original and the re-synthesized speech and then I'll show you on a more fine grade level what they look like. So these are spectrograms again, you can see they look very similar. Here's an original speech sound:

[Narrator 2] Okay. Answer me this. Why would anyone want to wear an overcoat in San Francisco, in the middle of summer? And here's the peaky speech that we re-synthesized.

[Narrator 2] Okay, answer me this, why would anyone want to wear an overcoat in San Francisco, in the middle of summer?

So they sound a lot alike and that was actually the goal. But if you look in close at the sort of vine structure of these waveforms, you can see the original, it's a bit kind of messier looking. And with re-synthesized, our peaky speech, you can see that the periodic part of the speech that gives it its pitch is very sort of well, peaky. And what this allows us to do is following each peak we can look and see, was there an ABR following that? We can essentially treat each peak as a stimulus and average all the responses that follow these peaks.

This approach works really well. And here's an examples of six subjects we can see nice wave fives on all of them. Here's the average across all the subjects from that study. Wave five was present in every subject. Wave one and three were often there, but not always and they were always small. Okay.

So what we can do to address the speed of the ABR is we can do this peaky speech. But what about the fact that ABR is small? And what do I mean by small? Well, here's an example of the ABR and I showed you sort of single trials before, but it turns out those were being nice. Here's what a single trial of ABR data looks like. So it's tiny, there's a lot of noise. And if you want to look across the entire auditory system, we need to also look at the auditory nerve, which is the very first part of the subcortical auditory system. And it's represented by wave one in the ABR. But that wave one, it's really small and it can be unreliable.

But it turns out there's another measurement we can make called the compound action potential, which also represents the auditory nerve activity. And you measure it with an electrode on the tympanic membrane, which is just a fancy word for the eardrum. And it looks like this. It is basically a deflection, a big negative deflection like this at just a couple milliseconds. So I mentioned putting an electrode on the eardrum, a standard eardrum electrode looks like this.

Although through some collaborators at the University of Utah, they have developed a much smaller electrode. And here's how they look actually on the eardrum. You can see the sort of this top one covers the whole eardrum, where this bottom one is much smaller and just touches a small part. And what's important about this new electrode is that it still gives very nice responses, but it's much more comfortable to have on for say a two hour experiment. And when we use our peaky speech and this electrode together, we can actually measure.

So here is, in red, we have the non-eardrum electrode that gives us a nice wave five. But in this same subject, we can also see that we got our compound action potential. So we have a really nice auditory nerve signal, which is again, that first part of the auditory system in addition to the later parts of the subcortical auditory system. So these special eardrum electrodes allow the signal to be recorded closer to the auditory nerve.

What about selective attention in subcortical areas? Another way of putting it is what are all those efferents for, all those downward projections that I showed you? They are commonly stated to be about attention or listening in noise. Those statements are often uncited. The published literature on this is super mixed.

We designed our experiments with speech to be able to test this idea. Can we present speech and compute the responses across the entire auditory system, auditory nerve, brainstem cortex. So our paradigm is going to look like this. We're going to present attended speech, unattended speech. We're going to record the EEG data and again, we're going to compute the attended response to those speech peaks in the attended talker. And the unattended response to the speech peaks in the

unattended talker. So I'll start with the cortex. And what's nice about this and nice to confirm is, we find really large effects in our cortex recordings.

And so here you can see, we've got zero milliseconds to 300 milliseconds. So this is our sort of slower cortical response. There's a very big difference that starts kind of early and by a hundred milliseconds, you see a very large difference that sort of goes all over the scalp. But now let's look at the auditory nerve and this is the one we get with our eardrum electrode and there's no difference. You can see here the attended response, which is in black, basically completely overlaps the unattended response. Excuse me, the attended response is in red, totally overlaps the unattended response, which is in black. Here we have the ABR, which is sort of the parts of the subcortical pathway that follow the auditory nerve.

And again, you see the attended and unattended responses totally overlap. We see no difference even though they were paying attention to the red one and ignoring the black one. We just see no effect. So to summarize our findings from this study, we see no evidence of a peripheral effect of attention. No evidence of an effect of attention in the brainstem. So we see no evidence of subcortical responses while still seeing the very large cortical effects of attention.

Okay. So what do we know? We know the auditory brainstem response, the ABR, can be used to assess subcortical sound encoding. And we know that these subcortical responses, we can measure them from natural stimuli like speech if we use the right stimuli, the right types of analysis, the right electrodes. But when we do that, these responses to competing speech recordings show that there's a cortical effect like lots of other studies, but no evidence of subcortical responses being influenced by selective attention. Some caveats though.

One is that there's an important word that I sort of glossed over, which is that EEG is only sensitive to neurons that are aligned. Spatially, they're sort of pointing the same direction. So neurons that look like this in the brain will make a measurable signal outside the head. But neurons like this, that kind of go every which way, it won't. And it could be that there's a subcortical area that is arranged like this kind of all willy-nilly that is being affected by attention. It's just invisible to us.

Second, even if selective attention doesn't affect subcortical responses, subcortical responses are definitely essential to selective attention. They're pulling out all sorts of information like the location of sounds and the timbre and the pitch. Lots of these things are related to subcortical processes, even if those subcortical processes aren't actually affected by attention.

So when does this process break down? Everything I've talked about so far and all the research, all the results I've showed you have been in listeners with normal hearing. But a lot of issues which span the auditory system can lead to trouble listening to speech when there's lots of noise.

Hearing loss and auditory nerve damage can affect the initial encoding of sounds. Hearing aids, which many people use, are some help but they amplify the speech you're trying to listen to and the noise. You get things like single-sided deafness and unilateral cochlear implant use, or some also bilateral cochlear implant use, can disrupt the spatial cues that are really important for attention.

People with tinnitus, even with normal hearing can affect central processes that lead to problems listening in noise. And then there are what are called central auditory processing disorders or CAPD, which are still poorly understood, but are thought to be cortical in origin. So there are lots of ways this really sort of complicated process can go wrong.

I want to take a second to acknowledge the people who were most involved in doing this research. Melissa Polonenko, who's at University of Minnesota. And then my students, Tom Stoll and Tong Shan, all of them have been involved in this process. And Tom in particular with the selective attention work that I showed you at the end. We have also lots of other people who have helped out and funders I would like to thank as well. And I thank you for attention and look forward to talking more now.

ANIL LALWANI - Dr. Maddox, that was really, really terrific talk. Lots of really interesting material there. We have a lot of interesting questions. Could we possibly just start by talking a little bit about

what the differences that you have seen in your studies between normal hearing versus hearing impaired individuals in terms of their cortical responses? Have you studied that or have others or?

ROSS MADDOX - Right. So that's a great question. So I have not myself, although I have intentions to, but others have. And it's a bit perplexing actually. So they do find differences in people with hearing loss versus people without hearing loss in the cortical envelope tracking.

And strangely, several studies have shown that the cortical envelope tracking is actually stronger in people with hearing loss than it is in normal hearing listeners. And there's ideas for why that would be the case, but it is kind of a perplexing finding. None of those studies though, just because the methods are fairly new, have looked at the subcortical processing, at least not in the way that we look at it. So that's something that I'm hoping to do sometime in the near future.

ANIL LALWANI - And Mitchell asks, it's sort of a similar question, is what are the changes with aging that you see in these responses?

ROSS MADDOX - Yeah. That's a great question. So I believe the results are similar, that older listeners also tend to show sort of enhanced envelope encoding, these bigger envelope TRFs. But again, subcortically, I have some guesses about what it might be, but can't really say yet for sure.

ANIL LALWANI - And do your findings in any way, any implications for how hearing aids might be designed or used or cochlear implants or both?

ROSS MADDOX - Yeah, yeah. That's a great question. So, absolutely. So with hearing aids, and a lot of this applies to cochlear implants as well, there's kind of two big things here. So the first is that what I was talking about today was just sort of let's get people in a lab, present some sounds, and look at the response to them. But there are also ways where we could potentially use these brain signals to inform hearing aids, what someone is trying to pay attention to.

Because with hearing aids, so the problem is not the physics or the math of being able to isolate one sound and a mixture of many others. If you get enough microphones, that's actually a fairly easy problem to solve. But what's really difficult about it is how does a hearing aid know which sound a person wants to listen to?

And so there's a lot of work right now aimed at basically decoding attention. And these futuristic hearing aids would measure the sound that's coming in and it would sort of separate it into sources and then it would also measure the listener's brain activity, essentially try to figure out which sounds a person was listening to and then amplify that sound for them. So that's one way a futuristic hearing aid might do this.

A second one is that what we're doing when we create these responses is we are creating a model. And by model, I mean essentially, a mathematical model, a black box basically where we are saying how does the brain take this sound that's coming in and turn it into this response? And one thing you could imagine doing, and there are several people including myself, sort of interested in going down this road, is there a way to take models like this, which we've been using in the lab, and train them on individual listeners.

And if you can model a person's auditory system and maybe model some of the things that are different, say in someone with hearing loss who's using a hearing aid from someone who doesn't rely on hearing aids, recognize those differences. And can we basically make a signal processing algorithm that is based on the model of that person, the individualized model of that person's auditory system.

The idea of sort of algorithms that are trained to help, to help enhance the neural responses for individuals based on their own recorded neural activity is something I'm really excited about.

ANIL LALWANI - One of the questions comes from another person. Well, Ross Maddox actually, is there any sense in which those with physical damage or deterioration to cochlear or other subcortical processes, can they try to compensate? That is, can you do exercises? Can you train yourself to be better at this attention hearing and ignoring the unintended hearing?

ROSS MADDOX - Yeah. So, yes. There's a whole area of research called perceptual learning, which is sort of about this idea. And I will admit, I am less aware of it than I really ought to be, but I know that there are some promising results there for helping people especially focus on particular spatial cues and trying to sort of improve their perception of auditory space, things like that. So there's definitely some promising work being done there. Although the experiments are long and difficult because they involve sort of bringing the same people back in again and again and again. So I have a lot of admiration for the folks who are doing that work.

ANIL LALWANI - I think, so along those lines, you talked about how you can re-synthesize speech, you call it the peaky speech. Can that be used for rehabilitation? I mean, can you take an original speech modified in some ways, so that being able to differentiate what you're interested in versus not interested with that? Is there any way to do that? Is there some work?

ROSS MADDOX - Yeah. So that's kind of the idea of these sort of personalized models that I was was talking about. And so the idea would be as once you understand how one person's brain converts sound into brain activity, can you then develop an algorithm? Maybe it looks something like peaky speech.

Although I'll say, we've tested people listening to peaky speech in noise and they don't show any improvement, but they also don't really show any worse perception. But that's kind of the hope is can we, once we understand the neural encoding of people with hearing loss, of different types of hearing issues, can we then design ways of tweaking the sound that they're getting that works with their brain to try and improve their perception?

ANIL LALWANI - Professor Anonymous Attendee is asking, have you considered studying music as a stimulus?

ROSS MADDOX - Oh, well, that's a fantastic question. Yes. And we actually had a paper come out one week ago, two weeks ago? Something like that where we did exactly that. So that was work that was led by Tong who was one of the people whose picture I put up there. And it's actually really interesting, it was difficult to study because the peaky speech, while it works for speech, it is based on a stimulus being speech, we sort of made some assumptions that way, but music doesn't work that way. So we actually had to develop completely different methods for doing that work than we did for this. And we eventually figured it out.

But it turns out that the way you do the analysis, and this is one of the things we talk about in the paper, has a really big effect on the results you find. So our first results showed that, oh my gosh, people's brains are encoding speech subcortically, totally differently when they listen to music versus when they listen to speech.

Then we developed a better way of analyzing our data and basically came to the exact opposite conclusion, which is that subcortically, the encoding of speech in music is almost identical. We did not find any big differences there. And then things start to vary more at the cortical level, which again is where sort for lack of a better word, the fancier processing is happening.

ANIL LALWANI - So along those lines, I know there's a question from Brian. He says, how do you know that differences between the cortical and subcortical responses is not just due to the difficulty of recording the subcortical responses?

ROSS MADDOX - Yeah. And so I mean, it is definitely difficult. What I would say is that before we got into the business of trying to look at attention subcortically, we first wanted to make sure that we could get good subcortical responses. And so we actually have a couple papers on just single talkers, people listening to single talkers and basically comparing in each person those speech ABRs to click ABRs recorded from the same person and showing that they're the same and that kind of thing.

And so, I would say the fact that we have gotten to the point where we can get a subcortical response to speech that so closely matches the click response, our sort of stereotyped ABR means we are getting good subcortical speech responses. And if you're getting good subcortical speech responses, then I think you can start to sort of make inferences once you get to this attention paradigm. But again, you can't prove the null hypothesis much to my chagrin.

But so it very well could be that there is something about our analyses or something else that is hiding an effect from us, an effect that is actually there. But until there is some new way of making measurements that doesn't currently exist, that works on, it would have to essentially work on different physical principles. It's about the best we can do or the best we can do right now.

ANIL LALWANI - Yeah. I wish you very much luck being able to do that, because this could answer a lot of our questions about our attention and inattention.

ROSS MADDOX - Well, it is. And I'll say, I would love to be proven wrong. I set out hoping to show the effect of attention subcortically. And as I pointed out, there are mixed findings on this. There are other folks who have done work and suggested there is an effect, but in our hands with our data and our analyses, which do give these really nice ABRs, we just don't see it.

ANIL LALWANI - So Katie says, she has two children with CAPD and her kids always say, I cannot hear anything if there are two voices at once. Yeah. Have you studied people who have CAPD?

ROSS MADDOX - No, but I would love to. I would really love to. I mean, especially because CAPD in particular, it's just very, it's really complicated. It's difficult to study, it's difficult to diagnose. And so getting these kind of responses. Again, starting at the beginning of auditory system all the way up, I think would be really illuminating. But again, we're not even quite to the point where we've published the results in sort of normal hearing adults, but there's all sorts of places, I would like to go with it.

ANIL LALWANI - We're coming down to our last couple of questions here as the time is nearly coming to the end. You talked about tinnitus near the end, an impact of tinnitus. Can you expand on that some more?

ROSS MADDOX - Not much, but I can try. So tinnitus again, it is very heavily researched. It's very tricky. It's very tricky to study. And one for the biggest reason is because, when we are giving people sounds to measure the responses to, we have access to those sounds, right. We decide what they are and we present them to the listener and we know exactly when they're presented, but we can't know exactly what people are experiencing with tinnitus. What we do know about it, is that it seems to be sort of central.

My current understanding is that it's not related to say the cochlea so much it is, as it is later in the brain. And it could be interacting with these processes in really sort of unpredictable ways. I guess is

about all I can say. It's similar to the CAPD and even people with hearing loss where the work just has to be done. But tinnitus is a really interesting one because even people with the same thresholds as, with normal hearing thresholds, people who experience tinnitus will tend to have more difficulty listening to speech in noise. So I would love to see how that's represented neurally.

ANIL LALWANI - So I think Kenneth is going to get the last question, which is just perfect for your talk. What is your advice for someone in a noisy environment? How do we function better auditorily?

ROSS MADDOX - Despite what it says on my name right there, I'm not a doctor. At least not a doctor like some doctors. But with that caveat, I think it actually really depends. So there's a lot of ways that the process can be disrupted and as I showed earlier in the slide and as Cherry suggested in 1953, there's a lot of things we do to solve this, this cocktail party problem.

And I would say to experiment, try focusing really on visual stimuli that if you have access to those. If not, you may find that you have a better ear and better ear doesn't actually always even mean that the sort of hearing thresholds are better in one ear. It's sort of a little more complicated than that, but you may find that that's helpful. I mean, the best thing to do.

Well, I would also recommend, and I've given this advice to many family members as well. If you're someone who could benefit from hearing aids, get hearing aids. They're not a silver bullet, but they are quite helpful. And yeah, failing all that, you can try to go to quieter restaurants over louder ones. Things when that's under your control, that's something when I was younger, I found myself at a lot of loud either because music or loud because they had all sorts of glass surfaces and stuff like that. But quieter places where the music is quieter and where there's less reverberation in the room can also be really, really helpful. So I would sort of take a Swiss army knife approach to consider all the things that could help and try to do as many of them as possible.

ANIL LALWANI - Well, thank you all for your attendance and Dr. Maddox for this informative presentation. We are so grateful to you, our community, for your support of our Emerging Research Grants program. Remember that you can donate to our efforts to advance better treatments and cures for hearing and balance conditions at hhf.org/donate. Thank you and please enjoy the rest of your day.