Speech Acoustics and Frequency Lowering
Recorded August 25, 2021

Presenter: Joshua M. Alexander, PhD

- [Christy] It is my pleasure to welcome back Dr. Joshua Alexander. He is an associate professor at Purdue University. Dr. Alexander's research is centered upon improving speech understanding and decreasing listening effort in hearing aid users. Welcome back Dr. Alexander, and I'll hand the mic over to you.

- Very good. Good afternoon everybody. And thank you for joining today's webinar. Topic of what I'm gonna talk about is obviously frequency lowering, but also it's gonna be a, quite a bit of speech acoustics embedded in this. We're gonna focus the talk around kina bookend it with clinical relevance and then kinda in between talk about three specific studies. Here is my disclosure statement. These are the learning outcomes, and I'll try to cover all of these in today's topic. Also keep in mind that if you are gonna be taking this course for CEUs on the bottom, near the slide number, whether the content is relevant to one of the questions that will be indicated by Q1, two, or three at the bottom here.

As you probably know, frequency lowering is ubiquitous, it's present in all of the major hearing aid manufacturers, a lot of their subsidiaries, they go by different feature names, they implement frequency lowering in very different ways. However, at the same hand, it's often one of the most misunderstood features of hearing aids. In part, because of a lack of information in terms of how the processing is being done, and unlike a lot of the other algorithms where we just have traditional noise reduction or directionality, which more or less operate in the similar ways. As we show it here in this list, the frequency lowering operates very differently depending upon the method, and that can have very significant outcomes for individuals who are listening to these different methods.

So I'm a quick do a survey kinda gauge my audience here. I'm gonna ask you three questions and that survey. will pop up pretty quickly here. First, I wanna know if everybody in the audience, do you regularly fit hearing aids? And if your answer is yes, I want you to answer, the second question is, do you regularly use frequency lowering and there you can use your own subjective gauge in terms of what is regular. I would say, maybe a couple of times a month, a certain proportion, 15, 20% of your fittings. And the last question that I'm interested in is on a scale of one to five, one to five, zero to five, rate your self-confidence in making decisions about frequency.

So if you're not confident at all, give it a zero. If you're very confident give it a five and I'll wait just a little bit to let these responses kinda trickle in here. So again, answer two, answer question two and three only if you're regularly fit hearing aids. Looks like you have to, in order to submit the survey, you have to answer all three questions. So we'll go from there. I participated, that's how I know. And can we see the poll results? Okay. So we got about one to five, 80% of our folks here regularly fit hearing aids. We regularly, those who regularly fit hearing aids, looks at about a third, regularly use frequency lowering, and it looks like a pretty even distribution in terms of the confidence scale, about it looks like around 25%, covering the first four categories ranging from not confident at all to fairly confident.

Good for you the 25% of the people that said you're fairly confident. Okay. That brings me to the next slide here, which is, I was anticipating there, wasn't gonna be a whole lot of high confidence in your ability when you fit the hearing aids with frequency lowering. And then a lot of that is because as clinicians we're trained to know before you click, rather than just making random adjustments in the fitting software, we wanna know what effect that's gonna have, acoustically okay, for our patients, and as I mentioned before, our algorithms like noise reduction, directionality feedback, suppression, you might think that we can use an analogy to ice cream. Those are all, it's like vanilla ice cream.

You have different varieties, different brands of vanilla ice cream, some tastes better than others. But if somebody tells you, you're gonna get vanilla ice cream, you know what you're gonna get. With frequency, lowering it's more like Neapolitan. You might get the chocolate might get the strawberry, might get the vanilla, and you're not necessarily sure what you're gonna get. We need to know what is happening underneath the hood. More specifically with frequency lowering than probably any of our other features that are available and hearing is because it actually changes the speech code. And given that we don't have a lot of information what's happening between the manufacturers it should not be a surprise that you probably not as confident as you can be otherwise.

Manufacturers, they tried to make the algorithms, your engagement with them as tangible as possible by using, maybe fuzzy label, subjective labels, comfort, clarity and stuff. We need to go beyond that. We need to know exactly, really what is happening for our patients in terms of the aided audio graph. So we need to explicitly know how sounds are remapped. We know we have at least two goals, one, the frequency lower sounds need to be audible. We need to know that what we're doing is gonna be able to feel heard by the patient. And then given that we're gonna wanna choose the, setting that's minimally invasive, that's gonna accomplish what we need to accomplish, but not cause for the detriment.

And the other reason why we might lack confidence is because, you really shouldn't, when we make a manipulation, okay, I'm experiencing difficulty communicating in noisy environments. Okay. So you might, you get a flavor for what sort of environments is just typical background noise, is it when you have multiple talkers or you're an engaged depth of directionality, or are you going to increase noise reduction? You a very specific goal in mind. You know when you've accomplished what it is that you've set out to accomplish if you've dressed patient. It may not necessarily be as clear with frequency lowering. And this is all we're gonna spend, Most of today's talk is on a number two here. Number one, I've already covered in a previous audiology online masterclass.

It's five hours total, but it's modular. So you can kinda pick and choose if you don't have that time. Obviously you get the CUSs is you've gotta do from start to finish. But if you just have a general interest, specifically with what's happening underneath the hood with individual manufacturers, there's that modular component to that. The second thing that I wanna talk about today, which is the potential outcomes, okay? So we really wanna have a very clear idea about what the potential outcomes are when associated with frequency lowering, because they set the stage for how we approach the fitting process. They affect the choices we make, when we select and program hearing aid, they also affect the method with which we verify, within the reusing realer measures or in the test box.

They also affect how we counsel the patients about proper expectations. And so here's the trick 'cause we've gotta find the right outcomes. And this is the burden, the onus isn't necessarily always on the clinicians. In some of this, the ball has to get thrown back in our court as researchers to develop these tools. And the point here is this is related to the Wikipedia, cause of the streetlight effect. You might be familiar with the drinker, who's looking for his keys under the street lights and the police officer comes up to him, What are you doing? Hey, I'm trying to find my keys. Oh, where did you lose them? All back over there in the corner. Well, he goes, why are you looking here?

Let's because we're that's, this is where the light is right? So it's an observation bias that occurs when you search for something where is easiest to look. So the tools that we typically have in our, tool box related to where we're evaluating, maybe speech and noise problems or audibility problems, okay? Things that are related to distorted cochlear processes. And the thing is, is what we're looking for is that, frequency lowering is not going to improve the distorted processing, that places along the cochlear, where we can already simulate with sound, okay? The problems that we need to focus on are problems that are generated with reduced bandwidth. And that's what I'm gonna talk about, mainly in this first study, what are the problems?

What are the benefits that we acquire as we increase the bandwidth of sound. To help us along that way, you may have seen this term, this term is maximum audible output frequency, affectionately known as the MAOS, has kind of taken off over the years. And basically what it is, this is a term that I defined many, many years ago. In fact, it was back in going through my records on my earliest use of this term was March 30th, 2009. Previous iterations of this concept known as known output aided output limit, and Darcia Jerking, the audiologist I was working with at the time at voice tone, national research hospital, he's like Josh, "You gotta make this tangible "for the clinicians.

" And we kept working and working. We finally settled on a maximum audible output frequency. And what it refers to is when you look at the aided speech spectrum, using your favorite hearing aid analyzer. So we have intensity on the left here on the Y axis, really our SPL perse frequency here, thresholds with the red line and really what we're referring to when we refer the maximum audible output frequency, without frequency lines is the first step is, how much real estate do we have to work with, okay? You max out gain before you get feedback, making sure the patients, the sounds are comfortably loud. What we're looking at somewhere between the average speech spectrum and the peaks, okay You have to use your judgment where you choose along this line, but in this case, we will choose the peaks of speech where they cross the threshold.

That's the maximum audible output frequency. And what it does is what it does for us, it tells us to at least two things. It tells us the severity of the deficit. So how much sounds are we missing? How much of the speech spectrum is in the inaudible region. It also gives us an idea of how much real estate we have to work with, meaning where we have to take some of that speech in the inaudible region and put it in the audible region. And of course the lower, this MAOF is, the more deficit you have and the less room you have to work with when you try to recode it. So the overall net benefit then from frequency lowering is limited by at least three following things.

The first is, is that when we have region of an audibility, okay, the particular method may have to distort this information, however, that might be, could be a compression and might have to truncate it. That is taken a limited frequency range of the total inaudibility range, taking a, not the full bandwidth per se. That's the first thing, okay? So we're not bringing all of that information that's available in the original input signal and putting it in the frequency loads. The second thing is, is that when we put that information, we have to put it somewhere. We have to put it where the patient can hear it. Oftentimes that is where there may be coexisting sound at the same time at those particular frequency regions.

So we run the potential, either masking the information that's already here or with our compression techniques, we actually have to move that information, in the audible region, that would be normally amplified through the hearing aid in order to come to the new frequency, lower information. The third thing then is that, and I represent this with the, up arrow indicating that we have to just by virtue of the fact that we have to provide amplification, for the frequency lowered signal, meaning that we are putting that newly recorded information on a place in the cochlea where there's probably some existing, outer and inner hair cell damage, albeit not to the same extent as it would be if we were to present it without the frequency lower.

Okay. So we're putting the information, we're putting a limited amount, perhaps putting a limited amount of information, in the inaudible region and where we're putting it, we're putting it at a place where there is also sensory neural hearing loss, okay? So these three things may limit the total benefit that we get compared to, if you just extend bandwidth. This is just a laundry list of the three things that I just described. And there's a fourth thing. The fourth thing is the focus of in that study, I'm gonna talk about in just a bit. This four thing is, is how much information, what's the size of the deficit. And that refers to looking at the average speech signal, where down here for normal hearing listeners, how much information is being lost when we reduce the bandwidth due to the severity of the hearing loss and the limitations in the transducers.

So in the sense four establishes, the number four here establishes the size of the deficit, that was the net loss of information due to limited bandwidth. And then one, two, and three influence how much of this deficit can be recovered. Just kind of get us up to speed in terms of really what sort of information in the speech signal. Hopefully this is a review for most of you. If we look at the sentence, children like strawberries, okay? So speech spectrogram, can we have frequency, the darkness of the bars correspond to the overall intensity. This is where we have the energy and the speech signal, speech rolls off at about 60, be proactive, roll off. Most of these off the shelf spectrograms will compensate by putting a pre emphasis filter plus 60 V.

To make it so that we can visually see it, let alone what our ears can hear, okay? When, especially when you have a hearing loss. Important point is if you look at the low frequencies, a lot of that low frequencies particular in the vowels, the approximates are characterized by frequency specific information that is formats. And where are these formats start and where they end relative to the sounds that preceded and follow them, heavily influence our overall perception of those sounds, up in the higher frequencies, different scenario. Most of the information is fairly broadband. It's noisy sounds, okay? Perhaps much more, better candidates for frequency lowering because it lacks that frequency specificity, it's fairly redundant across the high frequency bed.

Okay, so the first study that I'm talking about, we had 10 to 12 normal hearing listeners, depending upon we had, we tested constant perception, ball perception, 12 listeners for the consonant perception, 10 for the knowledge perception. The question is why would we use normal hearing listeners, we're interested in understanding more about what happens with frequency lowering and sensory neural hearing loss. And the thing is, is that first thing that we know normal hearing listeners, have in tactical processing along the entire length of the cochlea. And really what we wanna do is first one of our first question is, is we have a basic need to know what the effect of increased bandwidth is for speech information. When we have undistorted cochlear processing.

So idea is that if increasing the bandwidth over a certain range, doesn't improve perception for some sounds, we shouldn't expect that frequency lowering, when we lower that, that information, for those particular sounds, we shouldn't expect that, will we get improvements either if we don't see it normally. And then again, this idea is basic need to know what frequency lowering itself does to speech information independent of hearing loss. So the idea here is that if certain information is not available to normal hearing listeners, after frequency lowering, well, we shouldn't expect that it would be available to our hearing impaired listeners either. So this is kind of like a very early step in the process to understand what it is that we are looking at.

So can we tested consonant perception, we tested vowel perception. Consonant perception for each of the conditions, I'll talk about where we tested with 120 different nonsense syllables that are VCVS, vowel followed by consonant and followed by that same vowel. We had 20 different consonants, three vowel contexts, and I'll talk about the differences between those, A, E and U. We had two talkers, so not a wide sampling of talkers. So when we talk about male and female talkers they are specific to our talkers in this sample here. Vowels we tested they were 72 of them, a typical way of doing this is with some, a neutral phonemes at the beginning of the end, which we use the H and the D.

So he'd, hade, who'd had, hode, okay? Where the only thing that's varying in the middle is that vowel. We did present each of these with speech shaped noise, and that's important to know. In the case of consonants those 10 debut speech, sight noise and support speech, noise tends to be, speech itself tends to be, have a low frequency emphasis. So we're getting more masking with this noise in the low frequencies, meaning that the higher frequencies we might have actually have for our high frequency sounds, better signal to noise ratio in the high frequencies 'cause our noise was sloping. For the vowels, we had 12 vowels, we had six talkers 'cause we had fewer stimulus. So we were able to compensate with additional talkers, two males, adult males, two adult females, a boy and a girl.

What we're looking at here then is as a function of so we did this with low-pass filtering, or we can think of this as equivalently as our MAOF or maximum audio below the frequency. As we increase that along these discrete steps, what we're looking at is the improvement in consonant recognition. So we have it broken down by her male talkers and our female talkers, and the overall average here in black. These circles correspond to statistically, where we no longer see improvements in speech recognition with further increases in bandwidth. For these consonants, okay? Well for these talkers, sorry. The consonants are pretty it's 20 consonants, so we're sampling the range of consonants. For our male talkers, the Asamtotic performance levels about 41 Hertz, kilohertz.

And then for the female talkers, maybe shouldn't surprising that we do need more bandwidth, higher frequency input in order to achieve asymptotic performance levels. And that's indicated that's around 5,500 Hertz here. Also notice that because we did present these in background noise, we did present them randomly. So they couldn't acclimate to a particular talker or particular vowel context. Performance levels were less than a hundred percent. So meaning that no matter how much we, once we started looking at frequency lowering for these specific stimuli, we're not gonna, we're never gonna see a hundred percent performance, for overall 'cause we're limited by other factors. The other thing then too, is is that just notice of course, maybe not surprised that the females talker, until it's at a lower overall percentage than the male.

So clinically though, what we wanna think about is, is the maximum audio output frequency. How can we take this information and think about our clinical problems? This idea here is again, the maximum audible output frequency, this is gonna be determined by the hearing loss and the fitting itself, the coupling to the ear. And the idea here is that this is the would be then for these particular stimuli, the maximum performance improvement that you could expect, as you increase bandwidth. So when you have a lower cutoff at say, 2000 hertz, we could see that what, we have could get up to about 35% improvement, if we're able to extend that bandwidth all the way out to it's asymptomatic performance, okay?

So this kinda sets the limits then, how much could we get then with frequency lowers. This is showing you the amount of improvement as we increase the, as a function of the maximum model output frequency. So again, we're female talkers, we're still seeing improvement at some of these higher cutoff frequencies. That doesn't tell us the whole story though. We really have to start looking at these in the individual formings. Those are just the average and when really the proof of the pudding is when we start looking at the different classes of speech sounds. In particular so we see here, we have our Africans, which would be the Ja and the Cha. This is a very interesting case because when you have limited audiobility, they have a very reduced recognition rate, lower than all of the other categories, but very rapidly as you increase the bandwidth, performance jumps up to almost a hundred percent for these sounds like 95%, okay So what that would indicate then is, for those individuals, with the lowest maximum audible output frequency, the most of your hearing loss, if you give them just a little bit information, where should you expect to see the most improvement, would be those two sounds as Africans, the Cha and the Ja.

The other interesting thing here would be, let's say the fricatives, okay? So the F the, T-H the S the S-H the V in the Z, we see continued improvements, all the way up to bandwidth of 6,200 Hertz. And in fact, if we were just to parse out the S and the Z, that goes even higher. Pet's, Domecq, what's shown early that's about actually 20 years ago in the study, where in order to have optimal perception for male and female talkers, even normal hearing listeners need to bandwidth out to their highest when they studied, which was about, I think, nine kilohertz, for the. But as a class of sounds, we see that you can see further improvements with increasing the bandwidth for the fricatives and that's shown down here.

So we could still see improvements interestingly then, for our individuals who have maximum audible output frequencies of around 5,000 Hertz, which will be, you think about with our mild to moderate hearing listeners, we still see that there's still room for improvement by extending that bandwidth further, particularly for the fricatives, I guess the stops is another category where we do see improvements, to speech of the frequency lowering as well. Okay, and then I did promise that we would kinda break this down by the core articulated vowels. So we have of vowel contexts, interestingly, vowel Oz, when we have vowels consonants are preceded and followed by vowel A. You notice here that there are significantly easier to hear than when they're proceeded and followed by the vowels E and U.

In particular, let's look at the vowel E because we're gonna revisit this vowel context, in just a little bit. Notice the significant improvement for vowels or consonants that are core to collated with the vowel E. For low cutoff frequencies, low maximum model output frequencies, the performance for this vowel was actually significantly worse than for the vowel, consonants proceeding followed by vowel A and U. But as we increase the bandwidth, we see continued, we see more improvements, more room for improvement, the steeper slope here as we increase bandwidth. So it's those vowels, those constants, particularly those vowels, where we might expect to see the most improvement with frequency lowering. And part of the reason for that is the vowel E has a high second format, it's produced with the tongue, forward into the mouth that tends to push the, formats of the consonants that preceded and follow it.

It tends to push those into the higher frequency range as well. As well as because we have a narrower cavity then when we then go to co articulate, the frication and the following say S, is that also generates spectra that are higher in frequency compared to the other vowels. So that's part of the reason why vowel E, we see that shift with vowel E compared to the other two vowel context. Okay. Now here, what we wanna do, is just look at just what happens when we increase bandwidth for vowels. Really not a whole lot. And the reason for that is most of the energy in vowels is in the low frequency. So it's low frequency formats.

And what I've done here is I've just broken the vowels down by category. So this is basically by where there, how forward in the mouth the constriction, the articulation is for the different vowels, so we have our front vowels, our central vowels and our back vowels, okay? So the colors correspondents, what you see in the line here, again you notice that the front vowels, which are heavily dominated by the E and the A, we see the greatest improvement with increases bandwidth compared to the other vowel categories. And again, the reason for that is, again, if the front vowels tend to, they are associated with the higher second format frequency. Okay, so that's where the information is with these other vowels, the second format is already, is almost beyond the range of frequency lowering.

It's lower than what we are altering. So they're not really hearing much different. Whereas here we're actually giving them access to the second format where they didn't before. Also notice then for vowels that the asymptotic performance levels is much lower than for consonants, again because most of the information is in the low frequencies. So we're gonna come back to this vowel E because this is gonna indicate, Hey, vowel E, this is the one that has the most detriment of all the vowels will frequency lowering, but also is going to run the greatest risk for alteration vowel confusion, once we engage frequency lower again, because that, information for the vowel E and A and those other vowels are contained in that in the region where we might be lower.

We could also look at vowels by height, where the tongue height high in the mouth mid low. We really don't see again, if anything, the big take home here is the skin that, where we see asymptotic performance levels is much lower than for consonants. If anything, we see these, high vowels with a greater improvement with bandwidth. But I think that's probably because they dominated by that vowel E which is kinda carrying the rest of the category. Okay, onto frequency lowering. Okay, so that's our first study. Well, frequency lowering. So you might be wondering what is the information in the frequency lowered signal, and most of the representations of frequency lowering that you'll see, including a lot of the ones that you've maybe seen for me or that you have seen from me are shown in the frequency domain, right?

That makes sense because we're changing frequencies going into the hearing aid, but these representations in the frequency domain are not necessarily intuitive. It's hard to look at a spectrogram of the frequency lowering signal eyes. I see where the information is. So what is the information? So what I'm proposing is a different way of thinking about the problem, but in the time domain, and so that's what we have shown here. Here are three sentences, different sentences, that have been, low-pass filtered to simulate a complete loss of audiobility for the mid and high frequency sounds. This is the time wave form. And in the letters in red here, which correspond to a lot of those fricatives, stop consonants, you can just see them here.

These are where I've indicated, where we are losing information after the low-pass filter, right? And then the next slide I'm gonna show you is what happens when we use frequency lowering. So just let's look at, then what's in the boxes. So if you compare what's in the red boxes and you'd go back and forth, you can see what the patient gains in the time domain with the frequency lowered signal, okay? So unlike before they're actually able to hear something, whereas they're not able to hear it before. So what is the information then that potential information in the time domain signal here? There we go. There're couple of things. First compared to when there's no frequency lowering, the patient, so these are frequent, let me get back up.

These are time wave form, of these speech sounds after they've undergone frequency lowering, okay? So these are frequency, lowered signals. We have information that the patient's cued that something was said, okay. That as we have an indication about the presence of sound, the second thing is that we have more or less, we can tell from the APRA disady, of the time wave form, that the sound was likely from the fricative class, or it contains frication. The third thing is that compared to other sounds that have been lowered, we have information about their relative duration. So we have some duration differences across the different phonemes. So we get some information there. The fourth thing is we get some information about the relative intensity.

This guy's much more intense than this guy or that guy. Okay. So we get some information about from the relative intensity of the frequency lowered signal. We also get some information about the change in the tensity over time, that is the temporal envelope, okay? So those are the things I would think of. That's where perhaps we got to start thinking about where the information is and there's hopefully that this depiction is much more clear, what it is that a patient might be getting out of the frequency lowered signal. May be a little bit more intuitive than looking at frequency lowered spectrograms. All right. I'm gonna try to make this as simple as possible. So this is now looking into the second study.

The second study was the same group of listeners that participated in that first study. And we use their performance in that first study where increase bandwidth as a benchmark for how much information that they're getting from the frequency lowered signal. And the method that I used was nonlinear frequency compression. So you might have some familiarity with that. And the basic idea with nonlinear frequency compression, you divide the incoming spectrum, incoming sound into two parts. The lower part here in light blue, passes through unaltered. There is no frequency lowering. Everything in the darker blue here, starting here, this will all identify as our start frequency is then subjected to lower. And that's what I'm calling to the input compression bands.

So everything in the dark blue is going to be compressed. And where is it gonna be compressed? It's gonna be compressed into one of these four output compression bands. So the way that we interpret this as is that the narrowest compression band here, it's about 700 Hertz, we've got about almost 1400, about 2100 and about 2,800 Hertz. So we take all of that information and we squeeze it lot, or we squeeze it much less. And we're gonna do a parametric examination, so we're gonna four discreet output compression bands. Again, when we have a smaller compression band in the output, that means in order to get the same amount of information from the input to this narrow region, the output we're gonna have to do a lot more squeezing, compared to these conditions out here.

And we cross this then with six different start frequencies. So that gives us six times four different frequency lowering conditions. You'd never be able to do this in a traditional clinical study. Every single listener heard all 24 of these frequency lowering conditions randomly. So this really kinda gives us an opportunity then to really interrogate what it is, bottom up processing, that the patients are able to get out of the frequency lowered signal, okay? This might be a little bit confusing, but so hopefully this table might help a little bit if it doesn't that's okay. So this just kinda shows you the parametric manipulations, I had six different start frequencies, so you could start really low. You could start relatively high.

And then again, we crossed it with these four different compression bands. And so then again, this output compression band is if you add the compression band to the start frequency, that gives you, that tells you where the information is doing in the output. That is a maximum audible output frequency. That's what these entries here shows the maximum audible output frequency. So, okay if you start here and then you have a wider compression band, you're gonna have a maximum, a higher maximum audible output frequency. However clinically we think of that, the maximum audible frequency is a product of the patient's hearing loss and the fitting itself. It can't change it more or less. I mean, we try, we try to maximize it.

That's step one is maximize the bandwidth that you have to work with, okay? So really what we're faced as clinicians is working along the diagonal. This is what I am to work with. My patients giving me this amount of maximum audible output frequency. This is what I have to work with. Now we have four different candidates settings. We could start high, the frequency lowering, but it's gonna come at the cost of having more compression, or we could start low and spread that information out, over the audible range. In this particular study, not the third one that I'll talk about with our hearing impaired listeners. We do have a confound with the maximum input frequency. What this means is that as we go from left to right here, or as we go along this diagonal, we are bringing in more high frequency information.

Let's just keep that in mind. So if we look at those conditions that are those four conditions along the diagonal, where we kept the maximum audible output frequency fixed at about 4,100 Hertz. So we'll represent that with the star here. We can look at the four different possibilities there for this particular study. We can start low, okay? So again, all the information and the dark blue from here to here is going to be compressed into the rage and the red. We can start low and spread that information all the way up to the maximum audible output. The frequency that's gonna have a higher, output compression band, we're gonna have less compression, a lower compression ratio. That's what the CRS is a smaller compression ratio.

And as we increase the start frequency, keeping the maximum audible output fixed, that's what we have to work with in order to get the, information in, we're gonna have to increase the compression ratio, because we have less real estate to work with. And as you increase the start frequency higher and higher, your available real estate to work with in terms of putting that information becomes narrower and narrower. And if you wanna get the same amount of information in you have to do more and more compression. One way to evaluate how much information is in the full band signals transmitted during frequency lowering is this measure of efficiency. And I'll trying to keep this as simple as possible.

And the idea here is, as our comparison is the amount of information in the full band signal, the way we compute information, and we're gonna compare that to the frequency, lower condition. We quantify information analytically in terms of bits, so you don't need to know what bits are, but the way we get bits is by looking at the confusion matrix. So for each row as shown by the red letters, these are the sounds that were presented to the listener and then the columns, and is how they responded. It's confusious. Those along the diagonal indicate that are highlighted in yellow, that indicates that they responded, gave us the response that corresponded to the actual stimulus. These are the correct ones.

So in this particular condition, we sent 72 signals out or 72 of each phoneme, and then these would be the numbers that you would just do the math would give you the percent correct. What gives you information though? So it's not just percent correct.

The thing is this we think the values and the diag the off diag, where the incorrect responses are. And the idea there is that the more distributed, the incorrect responses are no, randomly are, that would be less information. Maybe the patient didn't hear it. Maybe they just couldn't process. They just couldn't get anything out of it. And they just collect either whatever was on the screen, or they said, okay, I know it sounds like one of those stop consonants, sometimes it sounds like this, sometimes it sounds like a T sometimes sometimes it sounds like a K, sometimes it sounds like a B okay, there's less information.

And so, as opposed to here, if you look at the S what S was presented, they only heard S 10 times out of the 72. 53, this is across listeners, 53 times, they heard it as an sh. This is the most common frequency error associated with frequency lowering by the way, is this hearing the sh during the SS and sh. All of the responses, most of the errors are concentrated in this one phony. That's actually indicate, this compared to some of these other conditions where the responses might be more distributed across one would actually carry more information, 'cause it says, okay, they heard something, but they just consistently got it wrong. And that's actually important to know, because it says, Hey, this would be something that we could more easily remedy with, maybe experience with training, maybe alternative signal processing to make these sounds, so they actually hear its not as the S but as not the SH but as the S.

So we basically take how much information in the frequency lowered signal, and we divide it by the total amount of information in the full wide band six. So just the important thing is when we see these numbers in the next two slides, three slides or whatever, it's not necessarily a measure of improvement. So where we see higher efficiency, it doesn't mean that frequency lowering is beneficial compared to conventional processing. And we're gonna use, we're gonna talk about that with our third study, okay? This just says of all of the potential, all the information, how much of it is actually transmitted with the frequency level. So the maximum you could have is well, not really actually have higher than maximum, we'll talk about that later.

But anything less than one would indicate that the frequency lowering can be improved and it's chosen is where there's room for improvement, whether that improvement be just with that particular subject, using training methods, using experience, or maybe it's time to go back to the drawing board and think about doing some alternative signal process or choosing a different frequency load. There's kind of a lot here. And the important thing is, is it's not, let's not get caught up in the numbers. So we have all of these different conditions. We have the 24 conditions. Let's just kinda take a big, big picture look at this. What this, everything indicated in green indicates these efficiencies are relatively close to one, meaning that after frequency lowering, we're pretty much getting all of the information with these conditions, that's available in the unlowered signal.

Remember we take, we're using that first study as our baseline. When we look at the maximum performance, in fact, it's the highest bandwidth study in those previous slides. That is the benchmark here. And this is basically saying with these frequency lower conditions, we are transmitting as much information as we possibly can. Notice that all of this, this is for the male talker by the way, all of these are with the higher start frequencies, for the male talker. For the lower start frequencies, we do see a pattern though, where as you decrease the compression ratio. So up here is we have a lot of compression, we have less compression. We actually do see an increase in the efficiency.

More information is being transmitted, and that makes sense as you spread it out, you're getting more information. With this particular study, you have to keep in mind though, as we're spread it out, what's happening though is we're spreading it out, but we're doing it at the high end, meaning we're pushing the maximum audible output frequency higher and higher. You don't have that clinically. Remember what we have clinically, the maximum audible output frequency is relatively fixed and we're working along the diagonal, 'cause that's what we have to work with. So the idea then is out of these conditions, as we go up the diagnosis, we increase the start frequency. And the look we are increasing the compression ratio.

We actually got, we are getting an increasing the efficiency. We're starting higher, we're also squishing it more, but we get more efficiency for the male talker, which kind of indicates kind of puts looks at, oh, the start frequency is probably pretty important to us 'cause we have a large amount of compression up here and we're actually getting more information. The other thing though, is it is confounded again, each study has its limitations about how you make the manipulations and the third study, the next one I'll talk about, we did a little bit differently, but here as we go from left to right or up the diagonal, we are increasing the overall amount of sound from the input that was available to the subject, so.

With our female talker, the scenario is a little bit different, okay? So females, remember they, their ASAM totic performance and baseline condition, it needed a higher bandwidth before we reached the maximum performance. So it may not be a surprise that there are fewer conditions here, particularly it's again, the high start frequencies, but where we have less compression. It's these subset of conditions where we see the efficiencies are close to one, that is we're bringing in as much information that's gonna be, that you're ever gonna get is available with these different conditions. And like what the male talkers as we saw before, as you, if you just go down the columns, as you decrease the amount of compression, you increase the efficiency, but you can't do that clinically 'cause in this case, we're just moving the maximum amount of frequency.

All right. So is a quick summary then of what we saw is a function of the males and females actually, if we average across all the entries in their head about the same amount of efficiency with all of the frequency lowering conditions. Remember our Africans the Cha and the Ja, they are the ones that saw this very rapid increase in recognition when we just increased the bandwidth from about 20 to, or sorry, 2100 Hertz to 3,300 Hertz. So no surprise here, the Africans had efficiencies, very close to one, for most of the frequency lower conditions. So these are gonna be your candidates where you think, okay, where am I gonna see the improvement this probably be for those individuals with more severe hearing loss, the lower maximum audio frequencies, you might expect to see the greatest improvement for those listeners, for those sounds.

Fricatives. Those are the ones that took the longest, if we looked at bandwidth, we needed the highest amount of sound in the input in order to get maximum performance. The efficiencies actually fricatives have the lowest efficiency, well 0.75, meaning there's room for improvement, whatever that improvement, however we might achieve that improvement, whether it can be through experience, through training, through alternative signal process, these are where we should concentrate our energy in terms of thinking about how we can improve the frequency lowering. And that's been actually the focus of a lot of my research and some of my patented algorithms. Approximants and stops are kind of in between. Nasals are dependent upon category. Once we kinda got beyond where we have the most, nasals are interesting that M and N , we opened up the nasal cavity.

We basically acts like a anti formate sucks out a chunk of that energy in the low frequencies, and that's actually Q. Once I kinda get beyond that where we're manipulations are higher than that, we don't see a lot or have close to maximum performance of below that they're fairly low. And then one of the more interesting findings then before move on to vowels is we have a correlation, moderate correlation between the efficiency and the response time, meaning that those values that were closer to one, the listeners were much more quickly to register their response. And then those that were lower, they took longer. That kinda makes sense, 'cause they're not getting as much information, a little bit slower response, and this is, separate from percent correct, right?

So the idea here is that to the extent that response time maps to underlying constructs like listening efforts, okay? Might be important because over extended listening time, that extra effort accumulates and might translate into fatigue, stress, less memory recall, et cetera. So again, another tool in our toolbox to start thinking about how we might quantify the differences between these frequency lowering algorithms. With vowels, we try to keep the simple, the ideas in a vowels is most of that information is in the first and second formats, and all of the frequency lowering of conditions that are, that were tested today were much, all of the operation was higher than that first format, and only in that second format, and really only for those front vowels, 'cause we're gonna revisit this in our next study here with the hearing impaired listeners, which starts in the next slide.

All of the efficiencies then are close to one, once we kinda get below a high start frequency and that's makes sense with what we saw when we went back to the baseline study, where we increased the bandwidth, we saw no further improvements, once it got beyond about 3,300 Hertz and very little actually going through 27 to 33. So most of the information needed from vowels is in the low frequencies, and that reflects here all of these efficiencies, close to one. If you notice the ones in orange, we're gonna come back to these, not these exact conditions, but this idea of very low star frequencies, high compression ratios, can be actually fairly detrimental to vowel recognition, particularly those front vowels.

And that's how we're gonna talk about this in the context of our last study with hearing impaired listeners. We had how many hearing impaired listeners, we had 16 approximately in each group. We had individuals with moderately severe hearing loss.

We artificially restricted their maximum audible output frequency at 3.3 kilohertz with low-pass filter, and so everybody had the same was listening to the same thing. But their hearing losses were moderate to severe for the mild to moderate group, we limited their use low-pass filtering to fix their maximum audible upper frequency at five kilo hertz. You could go to the paper when we tested vowels, just the same stimuli I just talked about in fact, but what I wanna focus on here are these high frequency phonemes.

These are the ones where we see some of the biggest improvements with frequency lowering. These are fricatives, but they're all proceeded by that vowel E. Remember that vowel E the high, second format frequency, there we saw the most improvement with our other baseline study, as we increase bandwidth, for constants that are core to collated them. So we ease, ease, ease, and ease, and they're all female talkers really pushing the envelope and high frequency where all the information is nigh frequencies. And when there is low speed, noise. So really packing all the information in the high frequencies. This first set of figures, then correspond of moderate to severely hearing impaired group, where we were limited the bandwidth at 3.3 kilohertz.

What's shown here in the dark bars with circled in red, this is what represent conventional hearing amplification. There was no frequency lowering. All they heard was the low pass filter signal, and everything went through amplification, compression, wide dynamic range, impressions. Everything was as like a clinical fit, okay? But they only had information up to 3.3 kilohertz, no frequency lower, this is ideal. Does frequency lowering improve perception? The answers for these phonemes? Yes. Statistically significant every single condition, we see suggestively significant improvement with frequency lowering for these stimuli. Real briefly then what we have are the light colored bars, the lower of the two start frequencies, we start at 1.6 kilohertz and then we had 2.2 kilohertz. We didn't have a lot of bandwidth to work with.

We didn't have a lot of real estate. So we're really kinda limited in our manipulations. Then within each of these three, these corresponded, how much of the input signal we brought in, okay? So if you're bringing in more signal, you're working with the same start and the same end you're bringing more in that means more compression, higher compression ratios. Notice what the low start frequencies the high compression ratios, we see statistically significant improvement. Now let's look at the S and the Z perception. There we see improvements of almost 50 percentage points. These frequency alarm conditions, particularly the two, that I just talked about. Low start frequency high compression ratio. 50 percentage points for these particular phonemes. So focus on these two with the asterisks here.

We're gonna come back to them on the next slide, 'cause they're gonna see what happens with vowel perceptions for these exact same conditions. Individual to mild to moderate hearing loss, okay. We had up to five kilohertz, okay. I'm lowered, three different start frequencies and two different compression ratios within each of those. Big story here is we do see significant improvements for a couple of these conditions with frequency lowering, compared to the conventional amplification, especially again for the S and the Z. Here didn't even matter which condition we see significant and pretty good, what did we get, almost 40 percentage points. Let's see the S we're closer 50 to 35 percentage points improvement with frequency learning. This is a group of patients that you would not normally think about having frequency lowering.

They had mild to moderate hearing loss. They don't need frequency lowering here. We see there is some, and that makes sense when we think back in that early graph, and look what the fricatives, we brought the 6,200 Hertz as a class up here, we see a significant improvements as we increase bandwidth. And then we said, we can go higher out, that goes higher out when we look at phoneme S for example. Needing more and more of the high frequency information to achieve maximum performance, and we can see that here in this results and you can see that you get those, when you engage frequency lowering. There is room for improvement, that improvement can be attained with these particular settings.

Oh, life is not necessarily, there is a trade off. I didn't say it. Let's focus on those two conditions. This is the moderate, severe group. These are vowel perception. What we have up here, these are of all the vowels, these are, I guess, the five where we saw any significant difference with frequency lowering. And unfortunately the difference is in the negative, it's a detriment, significantly worse with the frequency lowering. So these are kinda more the mid Mid Central vowels, front vowels I guess. And these are the two conditions I showed you before, where we had the most improvement for S perception. Low start frequencies, high compression ratios, and I kind of warned you when I said, look at the vowel efficiency, for those low start frequencies, those high compression ratios, where they took a nosedive.

We see that here with the hearing impaired listeners, that was what the normal hearing was. Now we see it here again, where we see significant detriments associated with those conditions. And the reason for that again is, our second format sitting right there at the sweet spot. So these are the conditions where we have 1.6 kilohertz frequency. Oh yeah and then 2.2 kilo hertz. But our second format sitting right there in that sweet spot. And that's what we can see here. So, compared to the control condition, what we see here is the spectrum, you could see where the second format lives, and then look at this high start frequency, high compression ratio. You see how much the second format has shifted.

And when we look at the vowel space, we can see where we have mapped out, here's the first format frequency. Again, going up to 1.2 kilohertz, that's all below where we're doing our frequency. See no shifting on the X axis. And the Y axis is the second format. Control is in the dark colored circle, and so what you can see here is with all is frequency, lowering conditions, particularly this one, we are seeing a significant shift in that second format frequency. And to the extent that the first and our second formats tell us pretty much all we need to know about all vowel identification, when we're, it shouldn't be a surprise that we see a significant increase in vowel perception, associated with these conditions.

That kinda brings us then to a close. This is kind of this idea then is what we see here is we're working on two ends. And this is just a hypothetical. We have this information transmitted on the, y-axis and fact, we can start doing that now when we talk about bits, we can talk about these being in common of values here, but this idea is for green, as we increase them on a frequency lowering, from left to right, we also increase, comes with increasing distortion, but it's ideas that you're gonna get to a point where you're gonna get maximum information transmitted for those phonemes. And as you do too much of a good thing, performance probably will drop, whether it be compression, whether it be putting it at the low start frequencies.

Vowel perception again, we never see an improvement with frequency lowering for vowels. We only see a detriment. And so it's this ideas that some of those same conditions, where we see an improvement for fricative perception, we see a detriment for vowel perception. And this idea is here is show taking a nosedive, because in fact, if I do frequency lowering at very, very low start frequencies and high compression ratios, everything's gonna go to muck. And the idea is that the overall average, then what we wanna achieve is with our process here, is we don't wanna do any harm. We don't want our patients to be, have worse performance when we turn on the feature, then when we turn it off and our overall idea then is trying to find that sweet spot when we look at the two together, where can, might not be the best, right?

Might not be those conditions where we can get the best S and Z perception, because it comes at too much of a cost for perception of some of the other vowels, but where can we get the best overall performance. We wanna give the patient more than we take. All right. So what do we have available to kinda in the clinic, these are actually research tools, I guess you can download them. They are different ways of evaluating high frequency, speech perception of high frequency phonemes, the UW plurals tests, which is looking at plural identification, cats versus cats, dogs versus dogs, phoneme perception tests, where they've basically, they filtered out certain portions of the signal, to give them to emphasize the high frequencies,.

The ORCA test from white X and the S-SH confusion tests I developed because some of the great, they're all six female talkers like we have about 120 words or actual words, and they differ only in their SNS H something that was, I developed that because I was interested in that confusion. How many times S is heard as the sh and how can we minimize that? So there's a subset that we thought would be good for a child. And with those would put into pictures, so you can see down here in those, the last reference sides, you could see all of that's publicly available for download. And then finally, just one caveat, everything I presented today, there's now some new ideas out there in terms of, everything I present is all the seems, all the sounds were coming for, they listened to them over headphones were coming from the same source origin.

And there's been a lot of great work by Bryan Monson at the University of Illinois, where he's examined, when we're in a typical cocktail party situation, we're facing the talker of interest and the CHF has extended high frequencies, and you get more of those then for other people who might be talking, but they're talking to somebody else, they're looking off direction. Those extended high frequencies are radiating away from the patient. And if you look at the spectrum of those two in this and those higher frequencies in a cocktail party situation, where we see an SNR advantage, a signal to noise and advantage for the talker of interest compared to the other talkers, and they're doing a lot of work examining hey, can people, can kids use this information to improve their overall speech perception?

Has not looked at it yet in the context of frequency lowering, but the society of that, what I presented today is, it's this very specific to those nonsense syllables. We don't have a lot of context. We're assuming this, the masking noise coming from the same place as the source. Let's bring this full circle, okay? So if you're not looking for, if you're not looking in the right places, you're probably not gonna find what you're looking for, and so what I mean now is we've looked at traditional outcome measures. This is what's easiest. This is low-hanging fruit. Those are things that we have available in the clinic, okay? Traditional pencil and paper tests. What does a patient report?

Measure word recognition. We probably aren't gonna see differences in outcomes associated with frequency lowering, using those traditional methods. We'd have to look behind us. We have to look at other tools as you saw the most improvement we get with frequency lowering and just what bandwidth extension is, for the phonemes S and Z, I say, Africans and stops depending upon hearing loss. If you have a pretty severe hearing loss that African perception can really skyrocket with just a little bit of extra information. Again, the stimuli presented, their limited context is all bottom up processing, or just fee CVS. Once you start throwing things into sentences, people can use contexts to fill in the missing information. So if you use extended sentences or passages, and there's not a lot of competition, you're probably not gonna see a lot of difference because people are filling in what they're they're missing, okay?

So this idea though, is that, where are we gonna see the improvements? Probably when we have the most limited context when the above information and the fricatives matters the most, and that's... If you go out on my other talks and it's been repeated elsewhere, the S for example, y it's one phoneme, but it has at least 26 different linguistic uses. And we use it about 8% of every single phoneme that we say, carries a lot of meaning, plural and present tense and so on and so forth. So when you miss that sound, you can miss quite a bit. So it's window, when that matters the most, when does it matter the most, when you have competitors, you have noise.

The information is very, it's not very predictable or it's new content, and you're just all bought more bottom up processing, or you don't have as much semantic knowledge. You can think of them as your patients, probably children who don't have as much to bring the situation top to bottom down, you're relying more on bottom up Again maybe we start seeing differences in listening effort and response time. We've actually kinda looked at these before, but now think about all of the works, we done on listening and tripolometry et cetera. And then as I just concluded before spatial processing, this is kind of a new, territory we might wanna start looking in is, what happens when you have a spatially distributed sound field, particularly with talkers who are radiating their eye frequencies towards talk of interest or the listener of interest, or away from.

Monday morning summary, what do you do with this information? I know it's a lot of the very hardcore, basic science kind of stuff, but what can you take out of this talk I hope, one is just realize the benefit from extending bandwidth of speech and normal hearing listeners and what we can then get from again, from frequency lowering and hearing a pair of limiters is going to be limited, right? It's gonna depend upon your starting point, that starting point as a maximum audible of frequency, what's the size of our deficit, how much real estate do we have to work with with recoding that information? What's your end point? How much of that input signal are you bringing in?

Are you bringing in a little are you bringing in a lot, if you're bringing in a lot, you might have to do a lot more signal processing or do more to the signal that could add or take away from the speech information. As we saw, which was kind of a running theme throughout, it depends upon your class of sounds. Depends upon the particular sounds. And that's gonna interact then as well as where, which frequency regions might you expect to see to improve. Again, focusing on those fricatives, the S and the

Z, stops are one where I do certainly see improvements with frequency lowering. I didn't show you that today, but the stops have a initial frication associated with them.

That little noise burst that can be, listeners can take advantage with the frequency lowering. And then the spectrum and the level, how high is the background noise, low frequency dominated. so you're forced to listen to high frequencies. Use sensitive speech material. Okay. I know that again, clinically, not all of this is available. And again, the onus isn't on the correlations of some of that are on us as the researchers to develop these, to normalize them, come up with some standardized norms. Inquire about ease of listening, plurality, female voices. This is where we see the biggest differences with frequency lowering. And then finally become more informed about the brands you use to better understand how different settings may help or hinder speech understanding, for that I'm just gonna refer you back to that masterclass that I referenced earlier in the webinar.

With that, I will gladly take any questions. I apologize for going a little bit over. I get carried away sometimes. Okay, in the Q and A okay. The question that I see in the Q and A, I only see one, okay. Given that the cochlear implant transmitted information is very severely compressed, can you speculate on the differential impact and vowel and consonant efficiency? I think the compression that we're thinking about with cochlear implants is not necessarily frequency, but amplitude, right? So you have a very narrow range with cochlear implants before you very narrow dynamic range before they could hear it before then, they experienced discomfort, which is separate from the frequency. And in fact, here's an interesting tidbit I'll kinda take, this another direction.

It's interesting because if you look at the fricative perception and perception of S as individuals with cochlear it was an early study done where individuals will cochlear implants actually do better that individuals with hearing aids on perception of those sounds. And if we wanna think of that in terms of efficiency, as the questions word here, those individuals extra cochlear implants might be achieve a higher efficiency for that particular class of sounds, because you have your state, you're able to stimulate those high frequencies with the cochlear implant, with the electrodes, albeit it there's a lot else happening at the low frequencies that might cause further problems. Well, thank you so much. I really appreciate your attention today. If you do think of something later, do let me know, contact me and I'll try to do my best to respond accordingly, thank you.