AI-Enabled Hearing Aids: What Emerging Technologies Mean for Your Patients

Presenter: Tammara Stender, AuD

You.

Now let's dive into today's course, AI Enabled hearing aids. What emerging technologies mean for your patients and Dr. Tammara by Dr. Tammara Stender, Director of Global Audiology for GN Hearing. Throughout this month, Audiology Online is showcasing groundbreaking advancements in the hearing industry with our 6th annual Industry Innovation Summit. Thank you for being here today. We encourage you to explore our course library for other compelling Summit classes.

A special thank you to our title sponsor, CareCredit, and all the innovative companies participating in the 2026 summit. It is now my distinct pleasure to turn the presentation over to our presenter, Dr. Tammara Stender. Thank you so much, Melissa. And thank you everyone for taking time out of your day to join me today to talk a little bit about something that we have been talking about for the last year. A lot.

It's very big buzzword AI and what it means in certain kinds of technologies, but especially hearing aids. Thank you again for having me. I'm thrilled to be here today. This is such a great topic because it is something that I think we are all rapidly learning a lot about and it's wonderful to be able to come together and chat about what it is, how to use it, and more importantly, how to make it help your patients have better outcomes.

So speaking of outcomes, here are the three learning outcomes for today and I won't go through them one by one, but I hope to satisfy each one of them. And if you ever have any questions, there will be question and answer. I will definitely look at it at the end of the presentation for sure. And then you're always free to reach out to me whenever you have a question as well as with any presentation, it won't be covering every single scenario. So there are some limitations and risks that are involved when any presenter tries to talk about a topic, especially one as broad as AI.

So let's just jump in, shall we? You know, every single time that I've been to any talk or presentation about AI, it kind of starts the same same way, doesn't it? What do you think of when you hear the term AI? And then all of these robotic kind of images pop into your head and about how it's wow, the future is here and it's going to replace perhaps us as hearing care professionals or it is something that we can't control and it has a mind of its own. So first and foremost, if we do nothing else today in this hour, I want to dispel any kind of a myth that that that is true.

I don't believe at all that AI is something to be feared or something to replace or something that will take over and control. If it is applied in a very intelligent way, it has amazing opportunities to make our lives easier, to make us more efficient, and to do things for our patients that we've struggled to do in the hearing aid industry up until now.

Artificial intelligence to most people means something like you are a passenger in a car that drives itself and it might be fun to look out the window and just gaze out and know that that car has got you and will take you safely to wherever you're going. But if you're a person like me and you're used to a world where you control things like I believe everyone here is on this, on this call, you want to have a modem of control. You want to have some kind of ability to steer where you are going and to have some power over the outcome you're going to achieve.

So instead of asking who's driving the car? Well, it's no one. The car is driving the car. Let's make sure that we keep ourselves driving the car in a way that makes us feel safe. If you think about artificial intelligence, perhaps part of the reason why we get so into the cautionary tales is perhaps the use of the word artificial.

Because we all know that things that are artificial is the opposite of natural. And when you're in healthcare, natural is usually something that you want to emulate as much as you possibly can. You don't want to artificially preserve food. You do not want to artificially give care. You want to give genuine, authentic, natural, human centered care.

So for the purposes of what we're talking about, I'd rather think of this in terms of augmented intelligence instead of artificial intelligence. And there is a distinction in the terminology. So what is it and how do you define it? Augmented intelligence is using technology to enhance abilities rather than to replace them them. Some examples are providing drivers with more feedback as they drive a car.

And in this image you can see that on the road surface there are some blue guidelines. And that is helping the car with its sensors identify if it crosses over a line. If you're drifting over into another lane, unbeknownst to you, it will make a beep inside the car. And that is something that a lot of cars have today. And it is firmly in the field of augmented intelligence because it's not going to prevent you from drifting over that line.

It might make a little buzz or it might make a ding or something like that to alert you, but it's not keeping you from doing it. It's augmenting your ability to know what's going on so that you can make these decisions. One of my favorite examples that happens almost all the time is on a road that I take very frequently. I am driving down the road and I have to swerve over the. Go around a biker or a pedestrian who's also walking in the road, where we don't have some sidewalks in our area.

And the car will alert me that I'm going over that midline and it's not happy about it. But I know because I'm driving the car that that's okay. So the car is giving me feedback to tell me that something is different than what the traditional way of driving would be. But I, as the driver, am making the informed decision that in this particular instance, that's what I want it to do. So in the end, I'm really driving the car.

Other examples that cars can do nowadays is alert you to invisible black ice on the road to alert you to a cyclist perhaps, who's not dressed very visibly to prevent you from falling asleep, to make sure that you are fit to drive. If it senses that your motions indicate that you might be getting a little, a little tired behind the wheel, it will vibrate the steering wheel, make a sound, etc, to try to wake you back up. So these are all perfect examples of augmented intelligence that are already out there that are not taking over or replacing the human, but just helping us be smarter and better drivers. It's not only happening on the roads, it's happening in healthcare and other areas as well, such as this scene of a surgery in which, if you notice, the surgeons are actually not, not touching the patient. They are controlling what happens with tools that are much more refined and small and specific than the human finger or a scalpel, etc, and so already this is being used in our industry, and it's not replacing them, it's augmenting their abilities.

And then, of course, we have smart watches that tell us if our heart rate's getting a little bit too fast. We have prosthetics, and of course, we have hearing aids. The one thing that I think we can all agree on is that the human brain is so marvelous in its capabilities and it is so quick and efficient. It is arguably the most powerful computational engine known. Ironically, this quote comes when you ask AI what is the most powerful computational engine?

AI believes that the human brain. ChatGPT believes that the human brain is the most powerful computational engine known. So if we're going to maximize the way that we process sounds in A hearing aid. Let's harness that most powerful computation engine, the brain.

Artificial intelligence, the way that it is defined generally talks a lot about the ability of the machine to do tasks that historically only humans do, such as reasoning, problem solving and decision making, but to do it in special ways that improve our efficiency. For example.

Actually, my slide seems to be a little stuck here.

There we go.

There's a huge field of learning called machine learning. And it is the ability of machines to automatically learn without being explicitly programmed. This has been used in hearing aids in the past. For example, if you use automatic dfs where you don't calibrate, there is a machine learning component to them where it will set up a digital feedback suppression that's a baseline and over time adapt itself to the characteristics of the patient's ear canal to eliminate feedback. So it has been in the field for a while, but we are expanding upon it.

And within that realm, that bigger realm of machine learning, we can talk about deep learning, which is a newer kind and a more advanced kind of machine learning. These are algorithms with brain like logical structures, and they call them neural networks because that is what they are called in the brain as well, that are capable of making decisions after being trained. So you can think of machine learning as it's trained to do something and it learns based on that training. Deep learning not only learns, but then goes one step further and applies these kinds of decisions in a way that is supposed to mimic the way that the human brain would do so. And this is what we really want to harness with intelligence, augmented or augmented intelligence, to be able to allow deep learning to make these decisions without overriding what the user wants or what we intend to have for our technology.

So what's a data driven deep neural network? It is the muscle behind deep learning. It's a layered structure connecting nodes or neurons. Every layer processes data bit by bit to recognize complex patterns. And if you think about it in a given day, in a given conversation, your brain does this all the time.

You're talking about one topic and then that segues into another topic. And at the end of the time, you realize that you kind of came full circle because this was actually expanding your viewpoint and your perspective. So that's what's happening here in artificial or augmented intelligence. The concept of it is not new. It's about 50 years old or 40 some years old.

But it has really never come into technology until the last few years in a way that's meaningful. One of the reasons for that is because it is pretty power consumptive it requires a lot of power and it requires a lot of data. And up until recently we didn't have the tools to be able to do that in miniaturized electronics. But be that as it may, it's an artificial neural network that mimics the human brain. And there are at least four layers.

There's an input layer where you input the stimulus when it's being trained. There's two layers. That's where it decides, well, if this, then that if, then statements and output layer where the output comes out. A deep neural network is trained on enormous data sets of real world sounds when it comes to hearing aid applications of deep neural networks. So that can be photos that have sound associated speech, etc.

It recognizes patterns like distinguishing a voice versus another sound that is not speech, like dishes clinking in the kitchen. So let's just take some examples here. These are just some fun examples to show where a deep neural network can make decisions that are appropriate and then ones where they fall apart and don't make appropriate decisions and how we can avoid those. So let's take an image. I'm in a suburb of Chicago.

So let's take the Chicago sun skyline on a beautiful day does not look like that right now. It is gray and there's snow everywhere and it's kind of miserable. But I mean if you like winter, it's not miserable. It is winter in Chicago, but let's pretend that it's not. And it's July and you see that there are boats in the harbor.

That is a beautiful sunny day. There's a light cloud layer. It's gorgeous. So what the computer sees, because it's not a human doesn't think that it's gorgeous. It just sees a bunch of pixels.

So in this video example you see that the computer is get is seeing pixels that are given a number response representing the color intensity. And this number of pixels can change slightly when the image change slightly as well. So in this case I blurred the photo a little bit to make it wavy. But as a human, I still see it as the Chicago skyline. I still know what it is even though it's not perfect.

In this pencil sketch version, it's very different than the first one because the colors are gone. It's much more rudimentary. But I still know that it is the Chicago skyline. That's human thought. That's my brain working.

So even though a change to the picture gives a new set of values, we must train the computer with many different samples of variations so it understands how to accurately Classify what city is this? Without over generalizing it or under generalizing it.

So the desired result in this case would be the input that would be the clear picture of the Chicago skyline. Hidden layers in that deep neural network where you have different characteristics, such as, if it is to be identified as the Chicago skyline, it must have certain characteristics for the buildings. It must have a body of water, because Chicago is the city on the lake by Lake Michigan. If it is a summer picture, there will undoubtedly be boats in that harbor and there will be pier characteristics for Navy Pier. And the statistical calculations that the system is trained on will assign a weight to each one of these characteristics.

So, for example, the building characteristics are very distinctive. They might be weighted at a point. 6. The fact that there are boats is not so distinctive. A lot of cities have harbors and bodies of water with boats.

So that's maybe a lesser characteristic, and that's weighted lower. Your desired result when you train the system like this is that it identifies as the Chicago skyline. The weights are optimized until the prediction is correct. So once this system is set up, the input is refed back into different variations to ensure that you get the same desired result identifying at the Chicago skyline, even if there's minor variations to the picture. So if you take this picture, what you want to have happen is with proper training, it will investigate the input in detail, apply those weightings to the characteristics and the hidden layers, and come up with the Chicago skyline answer, which is the desired result.

Humans are great at doing this. We can understand that things can vary in ways that are insignificant but not change the outcome. So even though this doesn't look like a color picture of Chicago, I see this as Chicago pretty easily. And wherever you are today in the world, I am sure you could be thinking of your city or your neighborhood and thinking of the ways that it would make sense to you based on the fact that you're a human and you've built these layers up in your brain automatically without having to put any effort into it, because that's the way the brain works.

So let's say that we don't train the neural networks correctly. And we start to say, you know what, it's fine, it'll be fine. Let's not train it enough. Let's just give some extra weight to the body of water. Let's give less weight to the building characteristics, as in this example.

What you're going to get here is you're going to get an error. The output of the system is going to say, well, there's a body of water. I see another city with a lot of body of water. It might be Sydney, Australia. And that's what it's doing.

Because there are suboptimal hidden layers and waitings to them. It's not identifying correctly. So it overgeneralized. It basically said this picture could be any city that has water by it with some buildings on it and a piercing. So let's call it Sydney, let's call it Dubai, let's call it any of these.

And that's not true, and that's an error. So in technology, we don't want this to happen. So these hidden layers have to be trained really, really well.

What if it called it this city? It also has a body of water. Not so much of a pier that I can see right here, but it also has buildings. So this is what we want to avoid in hearing aid technology. We don't want to train the system to make mistakes because even a human knows that Chicago does not equal Sydney, Australia.

Another thing that can go wrong is that if we add certain weights to these characteristics, we can come up with it being too restrictive and saying, you know what? I'm not exactly sure what this is. This is a pencil sketch of some city. But because it's not following exactly what it was trained to do, I don't know what it is. And then it is under generalized.

And what we get out of that is the terrible outcome of I have no idea what city this is. No result found. Searching the Internet and finding no result found or searching for something and not knowing what it is giving up essentially is also a terrible error to have in hearing aid technology. That would look like saying, you know, this, this situation is really interesting. There's a speaker and then there's noise, and then there's this coffee machine, and then all of a sudden it stops and it starts again.

It's too complicated. I give up. Let's just put us into some basic omnidirectional processing and call it a day. And that's not helping the patient. That's not removing noise.

That's not helping them listen to the environment or hear that speaker. So we don't want the system to give up or not find a result just when things aren't exactly the way that things are trained. So what we need to do is make sure that the training is accurate and we need to do that intelligently because as I mentioned before, this kind of technology takes a lot of power to run and it needs to be done thoughtfully and appropriately for humans. All over the globe. It needs to account for differences in languages, differences in tonality.

It needs to account for when I talk here in the middle, middle of my house when it's quiet, versus when I talk outside when there's traffic. I need it to be right no matter where I go in my life as a user. So the way that this is accomplished in resound hearing aids is with a two chip system. We are putting on the regular chip, which is the 360 chip, almost all of our features, that's our compressor, environmental classification, beamforming, directionality, wind noise reduction, all of that is on the 360 chip, just as it has been before we had augmented intelligence in our hearing aids. What is new is the integration of augmented intelligence in a dedicated chip that is designed to do exactly what it should.

Not overtrained, not under trained, but trained appropriately to help in many different situations where the user needs noise reduction. That's what it does, that's what it's designed to do. And we put it on its own chip so we can control the process better and make sure it's accurate and efficient. The reason why we put it on the other chip also is for efficiency, because we know that when we don't need it running, when there is not a lot of noise in the environment, then we don't have to send power through to use that chip. We can rely on the 360 chips.

So it's kind of like when you're trying to drive your car and you're not having to press on the gas because you're already going downhill. If the speed limit's 35 and you're going downhill and you don't need to press on the gas, then that's great. You maximize the efficiency of your gas mileage that way. So same thing here. We don't want to use power when it's not necessary.

So having this designated to do, specifically what it's supposed to do, helps us in terms of efficiency. And then let's not forget about the third and most important processor in our hearing aids, which is the human brain. We have to make sure that the user is in control of what they want to hear. That the kinds of noise that are being reduced by DNN in our hearing aids is not reducing things that the human wants to hear. And if there's more than one speaker in the environment, and a DNN is trained to pick up speech and make it more audible while reducing the noise, the DNN is not really going to know which speaker you want to listen to.

Because it was trained for human speech, but it wasn't trained for Sally over there and Michael over here. You, as the individual, make that decision. So we are making sure that the human brain and the individual wearing the hearing aids is part of the equation so that they can decide what they want to listen to. By controlling the environment, they are turning towards the speaker they want to listen to, which is a human habit anyway, when you want to maximize speech understanding and noise, and that is going to give you the best result in a difficult listening situation.

Like any other DNN system, we have an input layer, many hidden layers, and an output layer. So this is using this kind of technology, but applying it specifically in a way that it will help users with what is still the most common complaint, which is hearing speech in very difficult listening situations. We've come a long way and we are getting better. And this kind of technology sets us leaps ahead of what we had before we had it. So how did we train it?

Well, we know that training is very important. We don't want the outcome to be over or under generalized. So what we did is we used real life speech in noise sentences. We used sentences because words were not specific enough to be able to convey the kind of information that we wanted to, to avoid errors in, over or under generalization. Every language speaks in a sentence kind of form.

So this makes sense to do from a human perspective as well. And we want to ensure that wherever that person is, in whatever language they're speaking, it can be classified correctly by everyone. So the noise examples that were trained were things that we know don't have any speech qualities, such as traffic noise, airplane noise, babble, cutlery on plates, et cetera. And then these are some of the most important characteristics that we trained our system on. Frequency modulation, amplitude, spectral balance, waveform shape and time variation.

All of these are very good indicators of whether the signal is speech or whether it is noise. And all of them help make decisions that are correct. But we don't stop there. After the output is generated, we put it back through the system with slight variations to make sure that the weights are set correctly and reiterated in a way that's going to come up with the correct response.

So if you've seen some of our marketing materials, you'll recognize some of these numbers. As an audiologist, I don't necessarily look at these numbers as much as I look at for the outcomes that it gives to my patients, but I will share them here because it gives you an idea of the extensiveness of the training we trained our DNN chip on 13.5 million speech sentences. So that procedure that you just saw was done 13.5 million times. We worked across 3.9 million tune sound parameters, and we operate at 4.9 trillion operations per day. Fun numbers.

Great for Trivial Pursuit. I care about my patients, so let's move on to other ways we make sure that the training is accurate and efficient. One of the coolest things about DNN and how it's trained in the literature is that it doesn't always work in a linear past, present, future kind of tense. It actually is able to look at an object and then look at where it's been in the past to be able to determine what it should do with that object. So here's a fun example.

In this picture, this is a ball. Let's say, what direction is it moving? If you just look at it right now, I have no idea. It could be coming straight on my face. It could be going from left to right, it could be going from right to left, it could be flying away from me.

What if I looked back a few milliseconds before and saw that it was moving? That would give me more information about what to do with it. If I'm trying to get away from dodgeball hitting my head, then I'm going to move away from that direction, making a good decision in my human brain. But that's exactly what we do with DNN training as well. We not only look at an image as it's still photo, for example, we look at where it's been, we look at the sound and how it changed over the last few milliseconds to determine what we should do in the future.

So it's making informed decision. And what they call that in DNN literature is a recurrent neural network. It's an extra layer of sophistication to make more accurate and faster predictions about what to do based on what has already just happened. The human brain is great at this. We make a lot of decisions based on what just happened.

If I can see that my glass of water is wobbling, I can predict that it's probably going to fall. But I did that because I noted that it wasn't wobbling before and now it is. So using that kind of logic is really important in augmented intelligence too. The great thing about it is that when you take that extra step to use recurrent networks and you look back at what happened just a few milliseconds before, you're able to make faster predictions, and what you get out of that is a longer battery life. So a more power intensive method, the way of doing it instantaneously looks all the time.

It's scanning, scanning, scanning, scanning. It's like looking with a, with magnifying glass right at one moment in time and then the next moment in time and the next moment in time. That's taking a lot of power to do so. But if you look at it as a series of events and more holistically, you can conserve power by looking at the past and applying it to your future directions. That's really a key feature.

So this schematic shows how it all comes together in a hearing aid. As I mentioned, we've got the 360 chip with all the other features, and we have the DNN chip that's devoted not just to classifying the sounds, but also going one step further to separate out and not process the sounds that are processed as noise. This is called denoising. It is a term that is used in DNN terminology as well in other applications. But what you're seeing here in that final animation is that the speech is the green.

It's not being touched because it is being identified correctly as speech. That's wonderful. That's step one. Step two is the noise has been identified and there's noise reduction to it, but it's not being processed through the system because it is being identified as noise. So that reduces the noise greatly and really brings it down.

And this is a level of noise reduction that has not been matched by noise reduction that has existed before. Augmented intelligence. We've had noise tracker 2 for a long time, we really have, and we've made changes to it along the way. But this leap forward with DNN brings us so much farther into controlling noise than we ever had before because it's most, it's classifying and it's denoising removing the noise from that signal. So the output that the patient gets doesn't have that noise in it.

So since I just brought up our legacy noise reduction, which is called Noise Tracker 2, how does it compare? Well, Noise Tracker 2 worked by estimating the signature noise ratio per band. So in this schematic on top, you see the red is speech and the blue is background noise. There is a difference in modulation that really helps it identify that speech versus noise, but it's also frequency based as well. And what happened when you set up The Noise Tracker 2 is that depending on the setting or the environment, if you're using Environmental Optimizer for Noise Tracker as well, is that it will apply mild, moderate or strong Levels of noise reduction based on the signal to noise ratio and the overall level and other factors in the environment.

So it's a very good system. It is very good at estimating and reducing noise when it is detected based on the probability of it being in the environment. And that is based on our environmental classifier. The way it's been. How does it compare to the new augmented intelligence version of noise reduction that denoises?

We compared 20 participants with our old noise reduction or Our Noise Tracker 2 to our new noise reduction, which is what you see in intelligent focus in the in our hearing noise program. As you can see, it was around an mild to moderately sloping hearing loss. Everything was set up identically except for the type of noise reduction. So we could really see what the impact of the noise reduction was. We made recordings with hearing aids on a Kemar mannequin and then we used six sound seats and used female and male speech in the background of different kinds of noise, as you see right here, including a voice distractor.

Because we know that even though speech is something we want to hear, when there's multiple speakers, the other speakers you don't want to hear are distractors. And we want to make sure that you don't have to hear them if you don't want to. The recordings were binaural recordings and they were made at valid signal to noise ratios to emulate the way that real life conversations are. So we didn't contrive it to be, oh, this is a 95 DBA where it's going to be so noisy you wouldn't be able to talk anyways. We made it more like a real world busy restaurant level.

These are our results and we tested. In this case we were looking at the preference. What does the person who is driving the boat here? Of course they are the ones wearing the hearing aid and making the decisions. When do they prefer one kind of noise reduction over the other?

What we found overwhelmingly was a significant preference for the deep neural network noise reduction that is intelligent focus over noise tracker two in every single noise environment they encountered.

When you're thinking about how to apply this to the real world, we have to keep in mind that whole voice distractor. So you're having a conversation across the table at the Outback Steakhouse, but there's somebody talking really loud next to you. And your hearing aids were trained for identifying human speech. How do we make sure that we hear the person in front of us and not the person that's having a different conversation to the side in order to do this we have to marry beamforming directionality with DNN based noise reduction. The noise reduction is strong.

I mean that, that is a powerful system and it's going to remove the noise, but it's not going to always remove all of the speech. So in a real world situation like what I just described, you need directionality so you can control as the user who you are hearing and not have the hearing aid decide, well, that guy's talking louder than the person across from me, so I'm going to hear that instead. That's not what the user wants. That's an unintended result over generalization. We don't want that for our patients.

We want them to come in at their follow ups and say, yeah, it was great at the restaurant, I heard who I wanted to hear. I was looking at that person and I heard them great. It got rid of the other speakers. That's something that we want to make sure that our patients have. And it's not only about being able to hear well speech and noise, it's also about how hard you have to work to do it.

So this schematic is really cool. It comes out of a paper in 2012 done at the University of Oldenburg. And what it determined was that there's an inverse function for speech intelligibility and listening effort. So as your signature noise ratio, it goes from left to right on the x axis as it improves and it becomes easier and easier to hear. If we follow the speech intelligibility solid line, you see that at a certain signal to noise ratio you get a point where you're getting 100% correct because it's like a plus 15 signal to noise ratio.

People are not having a hard time hearing amidst the noise. They're doing great and they've maxed out. That's wonderful. Let's say that that does happen at, I don't know, say like about 8 DBs snr. But then look at the listening effort dotted line.

Of course, when the signal to noise ratio is poor, it's going to take a ton of effort, extreme effort to hear and they're trying really hard to understand and understand what's going on and repeat back sentences, et cetera. But there's a point where it crosses that signal to noise ratio S curve and then it maxes out, let's say around 8 or 9 DBSNR for speech intelligibility. But if you draw a line down to the dotted line at that level where it maxes out at signal to noise ratio for speech intelligibility you find that people are still having to put forth an effort. Yeah, they're getting them all right, but they're sweating, they're working hard. No one wants to work hard to understand because it causes fatigue.

And this is one of the reasons why noise reduction is used in hearing aids. The number one reason is to make things more comfortable and less fatiguing to listen to. So what we want to do is we want to get as far into the listening effort, being effortless as we possibly can, so that we can get not only 100% or very close to its speech intelligibility, but we also don't want people to be fatigued. At the end of the day. You don't want to come home from the Outback Steakhouse just exhausted because you were trying so hard to hear.

I'm just bringing that up only because I was recently at the Outback Steakhouse with my father who wears hearing aids. And it is loud and it is hard to hear. It is a fun place. But we, we switched him into the second program, which is the hearing noise program, with the intelligent focus enabled. And he did so much better because it was not only having him helping him understand my speech across the booth, but also making him relax.

He didn't have to work as hard because that noise had been removed. Much better than it could ever be with any other kind of noise reduction we've ever had. And much better than the first program too, because you know that you're locked in and you're able to really hear who you're talking to. And we had a great conversation, but that wouldn't have been possible without DNN enabled intelligent focus.

So with anything in health care, there's benefits and trade offs, right? Always, Always, always. So what do you get with DNN enabled noise reduction? You get increased listening comfort and noise. You get preferred sound quality, probably because it's more comfortable, and you also get reduced cognitive load.

So all of those are interrelated and they all fit neatly into the optimization of noise reduction. What do you have as a trade off? Well, it takes more processing power to apply noise reduction in a DNN fashion. Sometimes it increases your hearing aid size, and sometimes it can limit other feature combinations. So we want to minimize all these trade offs and we have strived to do that with the resound vivia.

But before I go straight into the sound vibia and what it does, I want to just summarize. A deep neural network is a data driven network that can be thought of as a machine's artificial brain but it is done so to be meaningfully trained to emulate what the real world is and how a human would react in that real world. What they want to hear, what they don't want to hear. And quality is so much more important than quantity. Even though we threw all of those billions and trillions things, I don't care.

All I care about is what my patient is getting and that they're getting it efficiently and in a way that makes sense and is easy for them to use. The DNN attenuates any sound source that's not speech because it has been trained to recognize speech. So no matter the background it will be attenuated. This is why it's very intelligent. And the patient benefits are increased listening comfort, improved sound quality, reduced cognitive load.

Combining it with directionality in a dedicated hearing aid program is really helpful because what that does is it really helps the user turn on that maximum hearing and noise help that they need when they're in a difficult situation. So intelligent augmented is a small but really important distinction from artificial intelligence. We don't want think of this as artificially. We don't want to replace anything. We want to assist the human in solving the problem which in our world is being able to hear with less effort in a difficult situation.

The goal is to augment what the hearing aid user already does already in their brain by being careful to not remove them completely from a noisy environment, but to make it easier to listen to that and to help highlight what they want to hear. And the combination of having a four microphone directional beam former and a DNN based noise reduction creates powerful help in noisy situations. This surpasses anything we've ever had. And it's really exciting because this is really where most companies are moving. This is where we are all looking at.

Wow, how can we use this kind of DNN to make it easier for patients, to make it easier for you so that you don't hear the same complaints over and over again that I did great, except for when I went to the Panera and then it was bad. We don't like that either. We hear it from time to time. We don't want to hear it. This is a huge leap forward.

But speaking of the clinic, how can you measure how much help you will a person will need in terms of hearing a noise? I bet you already know the answer to this hearing and noise testing. It's been around for a long time. The hint test, the quick send, all really good measures for determining more than what the audiogram can give you by itself hearing beeps or tones in quiet. So many patients have been like, well, yeah, I can hear that because it's quiet.

But what about when I get in noise? That's when I'm struggling in the real world. So the real world validity of hearing and noise testing cannot be understated. So for the purposes of this example, let's just focus on the quicksand test. And in that test, if you're not familiar, the user or the patient repeats back sentences in a background of noise that changes over time.

It increases over time to become more difficult. It's like one of those tests where it's like, you got it right and all of a sudden it gets harder. Well, it gets harder regardless of if you get it right in this test. But how it is scored is for each correct word, you subtract the total from 25.5 and this gives you the signal to no ratio loss. And of course, the reference for this is in our last slide here.

So an example of a score sheet for the quicksand test. These are all of the sentences they had to repeat. The words that are underlined are the ones that are scored. And let's say that and they are done at different signatures ratios. It starts at plus 25 signal to noise ratio, which is very easy, goes down to a 0 DBs in R, which is very difficult.

That's where the speech and the noise are exactly the same level. And you see how people do. You score the number of those underlined words correct, then you add them all up. So in this case, they got them all correct until things got a little bit hairy after the plus 10 signal to noise ratio. When it got a little bit worse than that, it got a little harder for them to understand all the words and repeat them.

That is your signals to noise ratio loss subtracted from 25.5, which gives you 5.5 dB. And then the guidelines of the quicksin indicate in the literature that if there is DBSNR ratings. So like, for example, this guy's 5.5 DB if he had done less than or equal to 3 DB. That's normal hearing in noise abilities. So for this person, if they come in and say, you know what?

I just. I'm. My budget's very tight. I don't know what I can afford here. But they've got this kind of a score on the quicksand.

You're like, oh, thank goodness, I don't. They might not need the highest technology to help them because they already do very well. And hearing and Noise, they're neurologically able to, their hearing loss is able to what have you. Each individual patient may have different abilities. You can just fit them appropriately with hearing aids and maybe it doesn't need to have all of the bells and whistles that would be out of their budget and keep them from actually saying yes to buying the hearing aid.

If they are like my example patient and they fall within the mild range of 3.5 to 7dB, this is where you say, okay, you know, sometimes it gets hard to hear in noise. And they're gonna be like, yes, of course. That's why I came, because it's hard to hear in noise. And that's when directional microphones and four microphone directional beamformers, that is using the directionality for each ear combined to make a beamformer more narrow response can make a lot of sense because that will help them understand better in noise. When they have a more moderate hearing ability in noise, then considering accessories is really important as well, because they are struggling more in a lot more situations.

And then if it's severe, consider all those solutions. So where would you put DNN noise reduction? Well, for a patient, the patients that I outlined here, I would start to think DNN based noise reduction would be a great idea when they fall between the 3.5 and up range, even with a mild impairment for hearing and noise. If you can get them some degree of DNN noise reduction, that will really make it more comfortable to listen in. So they might be repeating back the sentences, great.

But remember, on that S curve, you can get 100% speech intelligibility, but you were fighting to get it. No one wants to live like that. So DNN based noise reduction can be clinically relevant even for people with mild SNR loss, because it helps with the one thing that's really intangible and hard to measure in the clinic, and that's listening effort.

Just a quick note on directional microphones and what the differences are between two microphone and four microphone beamformers. Remember, all hearing aids that are directional have two microphones on them. So we're not talking about hearing aids with four microphones on them individually. No, that's not it. Using both hearing aids and a binaural pair to merge together the directional pattern and make an intelligent decision about where the speaker is to really focus in that beamformer to get rid of more of the noise outside of that speaker so that you can hear the speaker as best as you can.

The way that we've illustrated this before in our clinic, this was actually Done here in Illinois was we took our four microphone beamformer, which uses directionality on both sides and ear to ear communication to really be able to share that information and focus in on what, where the speech is and what you want to hear. We know that humans look in the direction that they are interested in. It is absolutely innate. It is a primal kind of ability. We know that.

So we're applying that because the human brain is a really important part of our processing. So we want to make sure that if there's a speaker in front of you that you want to listen to because you're looking at them, that you're hearing their speech in the midst of the background noise. So how you read this spectrogram is that here's the person, he's got a little pointy nose. We don't draw very well, but this is, this is an illustration. And he's looking straight up, that will be zero degrees.

And you see a, what looks like a record in front of him. And those are all of the frequencies. It's mapped from high frequencies to low frequencies from the center to the outer ring of the spectrogram. And the speech frequencies fall within the 1-5000 Hz range. In front of this user, we have a dark red area right around 5 to 7 degrees.

That's very close to being straight on at 0 degrees. That is where the beam former is picking up the greatest. That is where we're picking up speech sounds, because the majority of speech sounds lies in that region. But we also have directionality above 5000 Hz and that is more of a hypercardioid fixed directional response. And that is done purposefully to be able to give you relief from the noise that's in other directions, but also to keep you getting your spatial cues, things like interaural level cues.

Those are in those high frequencies. And if we put everything in a beam former in those high frequencies frequencies, that's really going to be harder for people to feel like they're listening naturally and not having the hearing aid zoom from one place to the next, which is disconcerting. So we add back that natural sound quality while still giving the signal to noise ratio benefit for hearing in front of you. And this was measured by having white noise and the hearing aid on the hearing on the mannequin and then rotating the head 360 degrees to generate this kind of response. This is cool because it allows us to look at what are we doing versus what else is out there.

And so here's like a diagram of the degrees where the beam former is happening, where the fixed directional is happening to create more of a natural listening environment and preserve localization cues like interaural level cues and what's happening at 90 degrees, at 135 degrees, where there's noise in the background and you don't want to listen to that.

A couple of other hearing aids on the market today also have beamforming directionality. And you can see from the top four microphone spectrograms that they are also focusing more towards the front. They don't do it exactly in the same way. Of course. Every company is going to be a little different.

We are very unique in that we have that delineation between the beam forming at from 1000 to 5000 Hz to the fixed directional response from the 5000 Hz to 10000 Hz. And we do that specifically. That makes us unique in the, in the market. And that's something that we're really proud of because we realize you want to hear well in speech and noise situations, but you also don't want to hear artificially and make it sound like you can't tell where sounds are coming from because you're losing too many localization cues, which is a danger with directionality. So we try to get the best of those worlds together in our beamforming strategy.

The two microphone directionality systems do have more of that red around the user. And what that means is it's picking up more of the sound, which was white noise here. And that's not necessarily what the user wants to hear all the way behind them when their microphones are set to directional. So although we, we measure this with white noise, what we're looking at is where is the microphone picking up. That's the most important thing.

And if you're in that mode, you should be picking up where you're looking at because that is going to help you hear better in front of you.

As a general rule of thumb, the narrower the beam, the higher your signal to noise ratio benefit becomes. So for 2 mics it's 3-4 db signal to noise ratio improvement. For 4 mics it can be up to 9, 10 db in some cases, depending on the setup in the lab. How much signals noise ratio improvement do deep neural networks provide for noise reduction? Recent studies across the board have indicated about 1dB of signals to noise ratio improvement.

Which is it's incremental, right? I mean, it's actually more like 13 to 15% in the real world is what it translates to. But what we really are focused on here is that if we're maximizing out our speech Intelligibility, if it's 100% on that s curve, we know we still have listening effort. People are having to work hard. And that's where we're happy to report that the DNN based noise reduction helps people with the intangible with being able to be more comfortable listening in noise without working so hard to do so.

So let's just finish here with a couple of case studies. I know that I'm kind of taking it right to the end, so I appreciate you all just hanging in here with me. Let's talk about Peter for a minute. He has a severe high frequency hearing loss. He goes to restaurants during peak hours.

Looks like he does that with his job because he is actually working as a restaurant consultant and he's having a hard time hearing his clients in noise. He's keeping up, but he's struggling. It's, it's a hard day's work. He's actually considering retirement due to this fatigue. His quicksand results indicated a 9 DB SNR loss.

Greta, on the other hand, has a severe loss. 18 years of experience. She lives with her family and cares for her two young grandchildren. She's always had hearing and noise issues, but recently she's even having more problems even with one to one interactions. And then Sam, he has severe hearing frequency hearing loss too.

But he's a new user and he's very active. He travels a lot. He's having a hard time understanding tour guides and navigating and airports because those airport systems aren't always the clearest and easiest to hear in the din of the airport. Airport noise. His quicksand results aren't so bad though.

He's doing okay. Four decibels. That's very solidly in the mild SNR loss range. And if we had more time, and I'm sorry that I'm running out of time, I would have us put some things in the chat about what kind of technology you would consider for each patient. And of course every patient's going to be different.

But for the sake of illustration, let's just go through here and just consider. What would you consider for directionality? 4 mic or 2 mic? How strong should the DNN be? Or should it be traditional noise reduction?

Are they mild enough that they don't necessarily need it? And if cost is an issue, because let's face it, the new technology gets priced at a higher price point. But if cost is an issue, I want my patient to walk out with something to help them, making sure that I help them and make their lives better, no matter where their budget allows them. What other settings would you consider? What level of technology suits it?

And what about streaming and wireless technologies?

So with Peter, here are some ideas about what came to mind if he were sitting across the table from me. With Peter and his nine DBs and our loss, I would recommend that beamforming directionality would help him because that's really going to help him focus on his clients and front of him and get rid of the background noise in the restaurant. But since he's being fatigued, I really want as much noise reduction help for him as I can get. So DNN noise reduction is really important for him. If he's at a price point where he can't afford that, then I would take Noise Tracker 2 and set it to a stronger level.

The remote microphone accessory could be helpful for him too because he could hear better the speaker in a really difficult outback situation. Outback restaurant or maybe some other streaming options. He probably is on his phone. So direct streaming from the phone, regardless of whether he's using Bluetooth, low energy audio or if he's using MFI or direct Asha streaming. And then if he's in public places, Auracast may be an option for him.

I would like him to have that option as well as perhaps a telecoil if he goes into those environments that equipped with that as well. Greta needs more help. We definitely are going to go with a 4 mic beamformer and if we can get her the DNN noise reduction to help her with those noisy situations, that would be optimal. But for her, streaming options are even going to be more important because remember, she had such a pretty significant SNR loss. She really struggles in difficult situations.

So talk to her like TV streamer, multi mic plus our cast. Where does she go and what does she need? Sam doesn't have as difficult situations, so he might be successful without having the highest of the high. But he will still appreciate form microphone technology because it does what we hope all technology does, which is support the human brain in the greatest capacity. So if that's an option for him, that would be what we would want.

Milder DNN might get him through the day. Traditional noise reduction might get him through the day if it's a cost issue. But either way our DNN is set up to be intelligent where it's not like oh wow, I wish I had turned off the DNN because that's too much. It's not going to be too much for anyone. What I'm saying is that it is designed to help your patients where they need the most help, even though everyone can use it.

It is really for these patients that want to listen more easily and with less fatigue and noise. And then of course, you can consider streaming options for him as well. I noticed that he was in an airport right now. The Frankfurt airport has been in the news. I know Frankfurt's in Germany, so it might not exactly impact me here in Chicago, but you can see this as a trend.

The Frankfurt airport is testing out its aura cast as we speak so that things can be less noisy. They don't have to come over the PA system with announcements of gate changes, ways to make it more pleasant of an experience for everyone. Because AuraCast will be enabled and people with AuraCast enabled earbuds and hearing aids will be able to access that information at a much higher signal to noise ratio than they ever had before.

So, in summary, what are we trying to do? No artificial intelligence. Let's talk about augmented intelligence or intelligence Augmented. You'll see it both ways. But we want to augment and help the machine, help the human, empower the human.

Make sure the human is in the driver's seat. They look at what they want to listen to. They can hear what they want to listen to without a lot of fatigue. We do not want to take over. We do not want to give the machine too much power and drive the human.

We want help from the machine to empower the human. We do not want it the other way around. And that's really, in a nutshell, how I would explain it to a patient as well. This technology uses what we typically call artificial intelligence, but it's used to augment your natural abilities so that you can still decide what you want to listen to without having the machine make decisions for you that you don't want.

And with that, let's use this technology to really help people in ways that we have never been able to help them before it was derived. Thank you so much. I have run right to the end, but I can, if you are able to stay on for Q and A. I'd be happy to take some questions at this point. Otherwise, I really appreciate you being here today and thank you for everything that you do.