Leveraging AI for Enhanced Sound and User Experience in Hearing Aids

Presenter: Joshua M. Alexander, PhD, CCC-A

Thank you for joining today's webinar Leveraging AI for Enhanced Sound and User Experience in Hearing Aids. It is my pleasure to introduce our presenter, Dr. Joshua Alexander. Dr. Alexander is an Associate professor in the Department of Speech, Language and Hearing Sciences at Purdue University, where he was recently appointed the Innovation and Entrepreneurship Fellow for the College of Health and Human Services.

He received his BS, Ms. and PhD in audiology from the University of Wisconsin, Madison, where he also recently researched speech perception as a postdoctoral scientist. Additionally, he completed clinical and postdoctoral fellowships at Boystown National Research hospital in Omaha, Nebraska. Dr. Alexander is dedicated to advancing speech understanding and reducing listening effort for hearing aid users.

He has developed a sophisticated hearing aid simulator that replicates the features of commercial hearing aids. This tool has been instrumental in at least 14 peer reviewed articles, five NIH grants, and the discovery of significant new knowledge resulting in four US patents. You can read more about Dr. Alexander on the course page for our website. Without further ado, I'm going to turn it over to Dr.

Joshua Alexander. Thank you very much for that nice introduction, for the opportunity to again share with you some of the information that I have to share just by way of disclosures, here's just the information that I have to disclose. There are inherent risks and limitations associated with this and all other presentations, so please review this at your convenience. In addition, these are the course learning outcomes that I'll try to achieve today. And just to make note to help you remember that there is a quiz that you need to pass at the end for CEUs and in the upper corner here on the slides where the relevant information is.

For each question you'll see a Q1, 2 3, 45 corresponding to those questions. All right, so artificial intelligence. Boy, that is everywhere. Everybody's heard of it. Of course, it's ubiquitous really.

And so it's becoming increasingly, increasingly prevalent in nearly every every aspect of technology, including hearing aids. And the widespread adoption of AI in our world then leads to important discussions about its role in audiology, in particular hearing aids, which is what we're going to talk about today. And depending upon your level of exposure to AI and hearing aids, some of this material may serve as a review. Other concepts may be new to you. However, just be mind that much of the information that's available to you and to me has likely come directly from the manufacturers as they have exclusive insight into how AI is integrated into their products.

I don't have any additional insights other than what is being publicly put out there. However, I'm going to take a broader presentation with today's presentation a broader perspective I'm going to try to provide an objective comparative overview of the technologies with AI that are currently available across the various hearing aid brands. Seems like every couple of months another brand is coming on board with a different variation of AI. And so it's important that while I'll be talking about a lot of AI technologies today, no single hearing aid includes all of the AI features that I will discuss, however. So throughout the presentation though, I will refer to specific AI implementations that are used by the manufacturers, but I'm not going to identify them by name to help maintain a focus on the technology itself and not the brand specific marketing that we see in our in our emails and in the product literature.

So with that, let's go ahead and get started. So over the years, AI and hearing aids has evolved significantly. It's not new. It has transitioned from simple automation from simple automation to sophisticated data driven intelligence. So it, while it might seem like it's new, it's formal use in audiology actually goes back at least 25 years.

And we can divide this progression from left to right into three major phases if you will, from rule based to data based to the deep learning with including cloud based AI. So the early, starting in about the early 2000s, this is more rule based. This means that the engineer has coded predefined sets of rules to handle various listening situations. So a simple example of this approach would be background noise or high apply aggressive noise reduction. However, these systems had a significant limitation because they didn't and couldn't adapt or learn from the behavior of the user.

Instead, they relied on the same fixed responses in all situations, regardless of how effective or appropriate those adjustments were for that particular user. And actually if you go way back, the earliest we go Back to the NAL NL1, the first nonlinear prescriptive method put out there by the national acoustics laboratory in 1999. So while it's not AI in the modern sense, it used an optimization based approach which can be considered a primitive form of AI. The A algorithm relied on the speech intelligibility index. Maybe you've heard that, maybe you've seen it, the SII shown some of our hearing aid analyzers just using that as the ground basis for what's the expected intelligibility with some modifications to what we use today.

And the whole idea here is to maximize intelligibility while ensuring that loudness remained at a comfortable level. So while this is important, it remained a static, it was non adaptive. I mean it would, it did what it needed to do it put out a rule based system for when we enter in an audiogram, what sort of gains should we apply at the various frequencies. Then as we kind of transition into the 2010s or so, then we start to kind of get into actually a little bit earlier than that. But up through the 2010s and a little beyond we get machine learning and database sort of systems.

So the idea here is we go from rule based to more data driven. And this shift was enabled by increased processing power and the availability of larger data sets that had come along in the early 2000s. And this has allowed then the application of machine learning algorithms to analyze real world usage patterns. One of the key advancements was the introduction of what we all know and love as the environmental classifier. And in fact there were some self tuning algorithms put out there that relied on machine learning.

And so now they can start to identify the hearing aids, could identify different types of sound environments, such as quiet rooms, restaurants, outdoor spaces, and then automatically adjust the settings accordingly. So unlike the rigid rule based systems, these were able to use real world data, help fine tune their responses and ultimately leading to something that's a little bit more adaptive, particularly from the user perspective. And then we kind of get into the late 2010s. This is where we start to get integration of user engagement with health monitoring. We'll kind of talk about this towards the end, particularly with this connectivity with smartphones, being able to monitor things like physical activity, fall detection, listening preferences.

And this is important because it kind of helps then allow for a more personalized experience for the user above and beyond just simple amplification. And then as we'll kind of get into is as well as we're starting to see AI being used in research and development, probably much more than ever before in terms of automated testing and verification. And then finally we get to cloud based systems and deep learning, which is kind of really what you're hearing a lot about now today. The idea here is that instead of relying on just the processing happening in the device, for example, with cloud based systems, we're able to access a larger scale of user data collected from a large group of users. For if you think about as crowdsourcing, if you will.

And so kind of with that let's go ahead and kind of jump into some of the differentiators. So let's think about then. It's always important then to kind of think about when we have new technology that's being put out there, you know, why, you know, what sort of benefits might they Have. What sort of limitations is this technology actually trying to overcome? So one of the primary limitations we think with traditional hearing aids is fixed noise reduction.

So traditional hearing aids, we know, often struggle in dynamic environments. There's a lot of movement of different sound sources, different sound sources coming on at different times, such as restaurants, crowded spaces, background noise is always changing. So older noise reduction methods applied a broad suppression to all background sounds without distinguishing between the unwanted noise and perhaps other important environmental cues that the user might want to hear. And so this kind of loses. Then a patient would lose availability to auditory information and create somewhat of a more unnatural listening experience.

Another challenge that was faced early on when we started to get to more options with hearing aids, is the need for manual adjustments. And this is no different now as we, as we get more options, there's much more that's available that could be the hearing aid can be doing. And it's not just not. It's not feasible for a patient to be doing this manually, to be making frequent volume changes, for example, or program changes. And that can be frustrating.

And particularly if the listening environment's unpredictable and the sound conditions are changing very rapidly. And so it would be just too burdensome to offload that requirement, manual requirement, and cognitive requirement to have the patient decide which all the options are available to go ahead and make that selection. We have limited personalization. So older hearing aid models focused on a more of a one size fits all solution, applying the same general sound processing for everybody. However, as we'll discuss, users vary in terms of what they want to get out of a particular listening situation.

And it's just not a simple matter of treating the hearing loss as it is to accommodate what it is that the patient wants to get out of a particular listening situation. And of course, speech and noise, that's always been a challenge. It still is a challenge. Point here is that with the application of AI, we can hopefully try to do a better job of separating speech and noise, reduce, if not anything, reduce overall listening effort and fatigue for the patient.

And so the idea here then is that as AI is being integrated into so many different aspects of hearing aids, it's kind of, again, it's essential to kind of take a few steps back, ask whether it's truly necessary for hearing aids, or if it's just simply another technological addition without much significant benefit. So to answer this question, we first must understand that making sound louder is not the primary challenge of hearing aids. I think hopefully we, we're beyond that. But, you know, if you ask the simple, if you ask a person, simply a layperson off the street, you know, what is it the hearing aids do? Oh, they make sounds loud.

However, as we know, volume alone doesn't solve the complex auditory challenges that users face in the everyday world. The most critical function of hearing aids, in fact, is control. To control the volume, they must selectively attenuate or enhance process sounds to ensure that the user hears what's the most important for that particular environment for them. And that considers it takes and then user multiple factors, frequencies that need to be adjusted for that environment, the timing of the sound changes, the direction from which sounds originate. And it's this level of precision and adaptability where AI can perhaps make a significant improvement.

So if we kind of think about, you know, basic signal processing as it is, you know, the earliest, easiest of which is dynamic range compression, where as we know, we adjust gain based on input loudness. There again here's another example of where it's just following a predefined rule. If the volume, if the input goes up, the gain goes down. Similarly, when we start thinking about traditional noise reduction methods and so the challenge, the challenge with traditional systems is that, you know, while we start increasing the number of variables, more variables are at our disposal for improving performance. We're still, you know, for from a human perspective in terms of trying to incorporate all that.

We're, we can only handle a finite amount of variables at one time. And to really get to know and understand their relationships between all of these different variables, and particularly when we get into more complex real world environments which are much more dynamic and much less static than maybe what we would simulate in a laboratory.

In addition, we have to consider that what's meaningful for one hearing aid user might not be relevant to another. And so just again, as we kind of get discussed, making this universal approach is just not going to adequately address individuals needs. So kind of moving on then just kind of again taking a big broad picture of AI because you're hearing a lot of terms thrown, you're hearing a lot of terms, you're hearing AI, you're hearing machine learning, you're hearing deep neural networks. So to a certain extent just kind of want to just take, you know, what's a big differentiator between these and just let's try to think about, you know, what, what is a. So AI just refers to the development of systems that can perform tasks that typically require human intelligence.

A basic dictionary definition. Important thing is that they learn from experience, make decisions based on data, and then ultimately recognize patterns to help identify trends and relationships. And so the key characteristics, though, is continuous improvement. So unlike traditional program programming, where systems follow a fixed set of rules, AI systems are designed to improve their performance over time by learning and process, by processing and then learning from the data. And this ability to adapt and refine the response is what separates AI from conventional technology.

So it does this by analyzing vast amounts of information. AI systems process large data sets. Usually the larger, the better. Quality is important from that data, then identifying patterns, trends, correlations, anomalies that a human observer cannot necessarily extract from the data using our traditional tools and methods. This is important because there's a lot of complexities involved as these various factors interact.

And then finally, as we discuss making predictions and adapting based on the learned information, applying the knowledge to make more informed decisions to improve accuracy and efficiency.

So again, AI is a big, broad field, encompasses different approaches. Again, when it comes to hearing technology, there are two primary types you're going to hear about. You hear these quite a bit thrown out. Machine learning, deep learning. They serve as the foundation for most of the AI that's in the hearing aids today.

The whole idea here again, help the hearing aids learn from the behavior, adapt to environments, and ultimately, if they can, improve their performance over time. Machine learning. Machine learning is a subset of AI that enables systems to learn from data, improve performance over time without being explicitly programmed. So instead of then following rigid predefined rules, these systems use algorithms to identify patterns within the large data sets. And once the system learns from the data, it can make predictions based on similar and future inputs.

So in the context of hearing aids, though, machine learning plays a key role in adaptive processing. For example, a hearing aid equipped with machine learning might analyze a user's previous adjustments that they do manually, such as volume, volume changes or adjustments to the noise reduction, and then apply similar settings automatically. The next time the user's in a similar listening environment, Then we have deep learning. Deep learning is a more advanced subtype of machine learning. It relies on neural networks with multiple layers.

And we'll get into this in the next couple of slides. Hence, we have the term deep. These multiple, multiple layers that's involved in the processing. However, unlike machine learning, these models can analyze massive data sets and extract hierarchical patterns, much like how our brains process information. And these networks then work by passing information through interconnected layers, and then as each layer as it goes through, further refines the information.

And this is probably best demonstrated with an analogy coming up. So, analogy for machine learning. So if we just think about an audiologist who follows a hearing Aid user throughout their daily life manually adjusting their hearing aid settings based on the environment. So at first, first the hearing aid, the audiologist relies on the user input to determine when to reduce background noise or make adjustments to amplification. However, over time, the audiologist will hopefully start recognizing patterns and understanding when and where the user pre certain settings.

Eventually, the audiologist becomes so familiar with the user's needs in particular environments that they begin making automatic adjustments without being prompted. And the longer this continues, the more better, the better the audiologist becomes at anticipating the user's preferences. So machine learning and hearing aid works of spiritual pretty similarly, you know. So initially the system gathers data based on the user interactions. For example, if a user is consistently lowering their the background noise in a restaurant, this information can then be noted as a preference.

And then the more data that's collected, the system can start learning the patterns, recognizing that in these particular noisy environments, the user prefers a particular setting of noise reduction or amplifications or directionality. And then once enough patterns have been learned, they can start making predictions automatically adjusting the settings before the user even needs to make a request to adjust those settings before they can get in and do that themselves. Deep neural networks. This is one, this is throwing out a lot DNN deep neural networks. These are often mentioned in discussions about AI because they're so powerful, I'd say it's worth taking a look at what do they actually do?

So a neural network functions like a team of audiologists, each with each team specializing in different aspect of the complex task adherence. So imagine a clinic where we have multiple teams of audiologists responsible for processing a patient's complaints. So each audiologist in a given layer here, so this is our layer, each as a layer represents a single neuron in the neural network. Well, all of the audiologists in this layer receive the same input from the previous team. So going from layer one to layer two, team one to team two, each individual audiologist will focus on a different aspect of the problem based on their own area of expertise.

And then as the data moves forward to the next layer, the next team refines and transforms the information to create a more structured and nuanced understanding of what it is that of the situation. It's an iterative process, okay, Going through and through again to help the system the neural network to extract the patterns, prioritize features and hopefully, hopefully at the end of the day make a better decision which is the output that you see at the end. So first stage, here's the input layer. This is where all the data gathering occurs. This first team then listens to the patient's concerns, gathers broad information about their difficulties.

They assess whether the user struggles in noisy environments, quiet settings, conversations with multiple speakers, so on and so forth. And once this general information is collected, then the data moves through these hidden layers that you don't see. So this is the input stage and every other layer subsequent to this is hidden. That's outside of the, of the operational process from view, if you will. But what's happening in these hidden layers, these are millions or billions deep in some of the larger systems that we are, what we hear about in the, in the popular press, maybe not so deep in hearing aids, but many, many layers deep.

And what happens is again, as they move from one to the next, they're going to analyze the specific details that they got. So each, again, each of these audiologists is like an individual neuron. They're each a unique combination of information from the team. So they're all getting a unique look at the information even though that they're in the same layer, okay. And they extract out what they need to do and then they pass that on to the next layer.

And as you go on refining the information and finally can tell you get to the output layer, this is where the decision is made based on all of this refined information that happened between the input and the output. So this is like our team, the last team of audiologists, compiling all the data to recommend the best hearing aid adjustment. For example, if this determine the users in a restaurant may automatically boost the speech while minimizing the background noise. Pretty basic thing. However, skills be it better than that, because one of the reasons these deep neural networks are so powerful is their ability to learn over time through a process called back propagation.

So if the system makes an incorrect adjustment at the output, such as removing too much background noise or making speech sound unnatural, it works backward to correct the issue. Each audiologist, now each individual neuron, each individual audiologist in this example will evaluate the information they passed on to the next layer and refinery their how they weighted the information from the previous layer to help improve the accuracy. And this iterative process, going back through, going through the network and then back through, learning from the decisions at the output, is what then helps it refine and become much more effective over time. The more examples it gets, the more feedback it gets in terms of the, the output, the better the system gets.

This is just a. I'm not going to go through this table point here is if you want to take a deeper exploration into some of the differences between machine learning and deep learning. Basic idea here is, is that if we kind of look at machine learning, machine learning is best for structured data, I.e. rule based systems such as adjusting settings based on user behavior. Deep learning, deep learning is better for more advanced sound processing, for example, to help separate speech from noise.

But there's a trade off. I mean, it's not so straightforward. So, you know, while deep learning is much more powerful, it comes at a, it's much more computationally expensive in terms of the processing time, in terms of the processing energy. And so to the extent that we see deep learning on hearing aids, it has to be managed very carefully and to a limited fashion, if you will.

Also in context then for what it means for a patient, the choice between machine learning, deep learning, it's going to depend upon the needs of the user, the complexity of the learning environment. Hearing aids that integrate deep learning, they're going to be at the forefront of the sound processing and trying to separate speech from noise. Machine learning based systems, they're, they're going to be better at, or probably most often used with implementing and practical solutions. For example, deter, you know, determining what sort of program switches to make what sort of environments the patient is in. And these are going to be much more practical because again, they are better for computational efficiency.

And then the current drain that's required to run these systems.

Okay, to the extent that these happen at different levels, again we'll kind of go through the supercomputers. This is more on the research and development side, devising these algorithms, putting them on the hearing aids. And so this, all of this is kind of happening before the scenes. We have cloud based systems. These are also supercomputers.

The interface really is kind of a smartphone that we're kind of thinking is interacting again with these, these supercomputers. It's not. And these don't necessarily operate in real time. Again, so the idea here is to learn from user input and then make informed decisions in terms of how it can modify future settings for the hearing aid. That's our smartphones.

So there's some, to a certain extent, there's some machine learning that's implemented on the smartphones. And the idea there is, we just got done discussing in terms of balancing this idea of computational complexity, the need for more sophisticated solutions against the more processing that's required, the higher, more battery that's required to run these systems. Some of these would be better off. They're offloaded onto the smartphones, not necessarily in real time, but as an Interface between the user and the hearing aid in particular environments. And then when we get to then the hearing aids themselves, there's maybe a couple of ways that we see AI implemented in hearing aids.

Again these are the ideas is that we have hard coded solutions that have been developed beforehand. And I'll have an example of one for a deep neural network noise reduction. The solutions from that are put onto the chip. And so the hearing aid is this then kind of running input and then spitting out output, but not necessarily adapting or learning over time. And then to a certain extent some functions we're starting to see more of this than actually being implemented.

They'll say deep neural network in the hearing aid, on the hearing aid chip itself. And the idea here is these are done at a much smaller scale. You can think of these as what I would call shallow deep neural networks. It's the same principle that we just discussed. Whether you have an input layer, multiple layers, hidden layers in an output, but they're just not so many hidden layers because every time you add another layer you're adding much more computational processing, much more demands for the current terrain.

So it requires a delicate balance. So to the extent we do see this, and there are some that have this deep neural network learning as they claim on the hearing aid, has to balance needs against for the user in particular situation as well as then you know, the complexity that's offered by these in terms of getting computational complexity and current.

So you know, what does, what do the, what do we have to, you know, get AIs everywhere. So AI and hearing aids shares similarities with other, you know, AI is used in other industries, things that we're familiar with, from streaming services to self driving cars. Easiest one is learning from data just like our Netflix or Spotify. You know, machine learning here is learned, is used to track what users watch, recommend settings accordingly. Just like how AI can be used in hearing aids to a certain extent to help recommend better settings for particular environments.

Real time processing just like we have real time processing and self driving cars where they have to instantly analyze the environment to detect obstacles, make split second decisions, rely on continuous real time data. So similarly some hearing aids have to with AI they process the environmental sounds instantaneously, analyze the inputs such as speech, background noise, reverberation in real time. Just like self driving cars must respond to a sun stop sign, hearing aid must immediately adapt to a noisy restaurant or a quiet situation. Pattern recognition and this is very fundamental to AI and being able to like for example cars distinguish stop signs, pedestrians, lane markings, same thing with hearing aids being able to differentiate between speech and noise, our hearing aids are a little bit more unique than what we our daily experiences. Other experiences with AI.

Okay, One of which is the hardware constraints. So unlike general purpose AI like our Siri or Alexa, worse out the process, standardized speech commands, hearing aids must tailor their solutions for each user in order to have the greatest benefit. So no two people experience hearing loss in the same way. And AI must learn that to adapt to each individual's unique needs in particular environments. Another distinction is the need for an ultra compact design and low power.

So self driving cars, they can have large scale computing, large energy sources that can allow them to do all the complex processing they need in real time. Contrast, hearing aids need to execute their processing decisions on a device no larger than a fingernail while at the same time preserving battery power and do all of the other things that they need to do that help the user experience.

What else do we have? Okay. Real time environmental classification. Okay. Unlike our voice assistants, Siri, Alexa, which operate in pretty static controlled environments, hearing aids on the other hand, must adapt to constantly changing soundscapes.

Whether in the quiet room or or noisy restaurant, walking outdoors, must instantly recognize what environment the patient is in to make those adjustments.

And then beyond all of this, you know, hearing aids, the impact, the impact of what it means for the daily user experience is, probably exceeds most of our other AI applications. So if our Siri or Alexa misunderstands a question, yeah, we might be annoyed and inconvenienced. However, to the extent these hearing aids make an incorrect decision, that can lead to communication difficulties and frustration and fatigue. So the importance of getting it right in terms of what it means for patients being able to more effectively listen effortlessly and apply their cognitive resources accordingly. So just kind of sum that up.

AI and hearing aids must balance user personalization, environmental complexity, real time processing within these small compact devices. And this is what makes hearing AI and hearing aids challenging. And it makes it unique to most of the other applications of AI we see in our everyday world. And with that being said, so here's a hat GPT. This is a, a fun acronym that I created just to kind of summarize the state of the world as it is right now for hearing aid technology guided by predictive techniques.

Not an official term, but the idea is just to kind of acknowledge the shift in how AI, machine learning, deep learning are being integrated into hearing aids. Being able to predict and adapt to users needs in real time. So unlike again traditional hearing aids that again are more rule based AI, the Idea is that these devices are more being able to apply more predictive techniques to optimize the sound processing. So ideally then for us as audiologists, this introduction of AI and hearing aids, it's exciting, it's challenging, but it does raise important questions about clinical decision making and about the role of AI in our overall patient care. And to the extent, we must question the extent to which these technologies really, truly improve patient outcomes.

So while they have the potential, they have a lot of promise to enhance speech clarity, reduce listening effort, the real world effectiveness remains, you know, in terms of, and it's going to depend, it's going to depend on data quality, its ability to continuously learn and refine through real world testing.

So it's kind of a cycle. So it's not a static system, it's just kind of an ideal setup. You know, ideally, you know, AI would be integrated throughout the life cycle of a hearing aid and refining, improving with future generations. So now, while no single hearing aid will fully implement this entire cycle, the ideal integration of AI is going to involve these four interconnected stages. Research and development, real time automation, user intervention and data analysis.

The first stage is R and D research and development. The idea is that these models are going to be trained on large amounts of data from a large number of users. The engineers are going to use deep learning, develop algorithms that are focused on, for example, enhancing speech clarity, reducing background noise. And this is then going to set the foundation for the next stage where AI is implemented in real world hearing aids. The idea here is that the automatics, this is where pre trained AI models developed from R and D come into play, or they analyze the environment in real time, make adjustments, things like noise reduction, directional beam forming, speech enhancement.

However, as we'll discuss, AI automatics is not perfect. Which brings us to the third stage, user intervention. While AI driven hearing aids, they really ideally would be adaptive for the individual users, get it right each and every time. The patients may sometimes still need to override the settings. Okay, when the needs don't, when the settings by the decisions made by the hearing aid or the automatics don't fully meet their needs.

And if the user's consistently making these manual adjustments in particular environments, then AI can learn from these preferences and then apply them proactively in the future. And that brings us to this last stage analysis. So particularly these hearing aids integrated with cloud based systems, being able to, that can then analyze the user interactions, how they interacted with their computer, their hearing aid, in particular listening environments and analyze that into patterns and then feed that information back into the research and developments so that you can future updates or future products can then improve upon the existing features inside the hearing aids. And that then brings us to R D. So again, one of the most important applications of AI and R D is algorithm development.

So this occurs long before the hearing aid ever reaches the users. Okay. And then these may not then require any direct interaction with the user. So the idea is, we discussed, the researchers are training the AI models, massive data sets and the tens of thousands into the tens of millions to expose them to a variety of real world examples, real world environments. And this is where the quality of the data set comes into play.

The larger and the better quality, more representative these data sets are going to allow the AI, these algorithms are developed at this stage to perform better, to be able to recognize more complex patterns and to optimize the hearing aid settings on conditions that it was never trained on, done in an ideal world. It also enables an adaptive algorithm that can adjust the hearing aid settings, as we discussed, automatically based on the incoming environment, rather than as we discussed the earlier state of affairs where they had to rely on fixed pre programmed rules to make these adjustments in real time and learn from the user interaction. Another interesting application is predicting speech perception, which helps the output, which speeds up the algorithm testing. So some manufacturers use AI driven speech detect systems as their standard, by which the predicting how well the hearing aid, a typical hearing aid user will benefit from the particular solution. And they use that then to train up the algorithm based on that algorithm output.

So basically using speech detect systems that we're familiar with from Google, for example, to predict how a human listener will perceive a sound in a particular test condition. And that can then instead of then requiring hundreds of participants for the initial trials they can get, they can refine their algorithms before they even bring in the first participants for trial testing.

All right, so let's take our first application here. So, traditional noise reductions reduction in hearing aids long relied on these rule based systems. We had discussed things that engineers have conceived of in terms of if this, then that. However, the important thing is as we know is they over from lots of experiences dealing with our patients that they do not effectively have not been effective in minimizing noise in complex environments. This is where deep neural networks can offer a more advanced, a more data driven approach by learning from real world noisy examples.

Let's take an example from this particular application. The process of training deep neural network noise reduction models begins with capturing real world noisy environments. In this particular implementation, the manufacturers used a spherical microphone Array to record diverse soundscapes such as subway stations, restaurants, outdoor settings. And these recordings preserve the three dimensional spatial characteristics of the real world noise. So there's, I think there was 32 of these microphones.

And next these noisy environments were reproduced in a controlled lab setting. So a loudspeaker array ranged in an anechote chamber to recreate these soundscapes as they were recorded. And then putting that same sphere, that same microphone array right in the middle, then to extract the properties between the original environments and the, and the unique transfer, transfer function, the new transfer function from each of these microphones so that it could basically invert the process and then recreate these soundscapes in the anechoic chamber as if they were in that original environment. And that's what we see here. This is actually not a person.

This is actually to simulate a mannequin. Okay. So we put the mannequin in the simulated noise environment. The hearing aids are placed on the mannequin at the center of the loudspeaker array. And then we expose the mannequin to all of the recreated background noises that were recorded earlier.

And then the hearing aids do, they're programmed as they are to capture their everyday processing of the gain modifications for particular hearing losses. And then as we know with these mannequins, there's microphones inside or the ear canals would be, and record the output of the hearing aids in response to these particular sounds. And then we get to this very important stage which is to develop a much more realistic training set for a neural network. The model was trained on individuals wearing open back hearing aids. Now this is a real person listening over headphones to the, to the recordings.

And what they're doing is they're now, and this is the idea is that it's as if it's in their head because they're hearing this and they're un occluded, which is important. And they're adjusting their vocal effort accordingly. So it's as if they are in that loud environment. So as opposed to some of these other variations where you get a lot of noisy recordings and they get a lot of clean recordings in a quiet environment. Well, that's not how we normally speak.

We normally adjust our vocal effort or frequency emphasis when we are exposed to background noise. And that's what they're trying to do here is recreate how people actually talk in these environments to get a much more realistic and a much more accurate sense for the incoming acoustics when they reach the microphone for the patient. And so that's what's done here. So now what the important thing here is in a stunt over headphones because then they have a clean recording of the of the speech. So the speech in quiet, but as if the patient were individual were talking in the noisy environment.

And then they have the recordings of from that that environment.

And then they have then the two simple samples, they have the clean speech recorded here and they can mix it at diff. Mix it accordingly with the noisy speech. So you have a clean speech sample and then you have a noisy speech sample and then they can you then they use the, the deep neural networks, they use those two samples to basically say, hey, what needs to happen in these hidden layers to make the output, the the acoustic output as similar to the clean speech as possible? How can we make the noisy speech look as look is similar to the clean speech? By modifying all of the intermediary hidden networks and then apply that solution to the hearing aids.

The solution to this process then applied to hearing aids to have the deep neural network solution in the devices themselves. And the advantage of that is that so in this particular case, the input to the algorithm in real time is using the information from beamforming to estimate the signal to get a better estimate of the signal to noise ratio, individual time and frequencies.

And so the interesting thing is that behind the scenes, the hearing aid can be using the beamforming information while presenting to the patient an omnidirectional signal. And the idea here is that this can then allow, ideally allow the patient to benefit from the directional information that is the better estimate of signal to noise while presenting with a more natural signal. Because if we just put the patient in beamformings, we know that has a trade off because we're going to isolate them from the situation from sounds coming from specific directions. And that comes again at limitations. So a desired speaker moves mid conversation.

Okay. Then they may have to play catch up to try to figure out where the intended speech is. Or different talkers jump in into. And if it's a very dynamic situation, that could cause them to miss some sound shifts. So to a certain extent there are many environments where we want to have the patient opened up and so present them with an omnidirectional signal.

But at the same time, in an ideal situation allow the noise reduction to have the same effectiveness at improving the signal to noise ratio as our beam forming.

And then that kind of then brings us back to the automatics.

So one of the key challenges as we know is as faced by hearing aid users is the need to navigate different acoustic environments in their daily lives. So a person may move from a quiet living room to a busy street, to a noisy restaurant, back to the park. All of these environments, in a short amount of time, presents unique listening challenges and requires different settings to optimize what that patient's trying to get out of it in terms of speech understanding and listening comfort. So without automatic classifiers, patients would have to manually adjust their settings to accommodate each and every new listening situation. As we discussed earlier, that's a significant physical and cognitive burden that most patients aren't going to do.

So that's why very early on we see the environmental classifiers and decision making going way back at least 20 years. The point here is that can we do a better job with our classifiers using artificial intelligence? So quick thing is we just kind of take a view. What do classifiers actually do? Okay, they don't just make a single decision as kind of saw in the spinning wheel here, they assign probabilities as shown by the area on the circle here, to different listening scenarios such as quiet noisy environments, speech, noise, music.

And then the hearing aid selects the program with the highest probability and then we'll dynamically adjust its settings for that particular environment. And so traditional hearing aids have often, as we've discussed, relied on rule based classifications. So which is if this, if this, if this, then that logic for them in this environment then do this that doesn't have the adaptability that is often needed in complex real world weight listening conditions. And so the idea here is that these AI going back again fairly, you know, using machine learning, okay, to train these classifiers on large data sets. And as, again, as we get better with computational complexity and efficiency, the R and D side have even larger data sets and more complex interactions to get a better representation of the diverse listening environments that a patient might find themselves in.

And again, to offer them flexibility and then adaptability to enhance the overall user experience. And, and then finally, just as we just indicated, AI is more than is more, is not used just for classifications. Also then once it determines what environment's at, then determines the optimal solution, how should it respond to optimized the settings for that particular patient. Okay, so we'll kind of go through that. We've discussed this.

So point here is that then our classifiers can and do fail the acoustic classifiers using the information gathered from the microphones. So they're not without their flaws. Even the best assist best systems, even the best systems out there today sometimes fail. And that can often lead to, for example, they could lead to a misclassification or incorrect adjustments coming out of that. That.

And so one of the main challenges is to identify the listening environment. So again, they have to determine what sort of environment the patient is in. And studies have shown time and time again that identical listening environments can be classified differently by the hearing aids from different manufacturers. And this then demonstrates the fact that they're seeing some inconsistency in their logic in terms of how they're classifying the environment. If they're all arriving at a different, different classification, the easy ones, speech and noise, or speech and quiet, sorry, highly noise situations, those are fairly easy.

But the dynamic environments where there's a lot more going on in the acoustic scene, that's where you start to see the big discrepancies between the manufacturers. So we don't know who necessarily who gets it right, but we always know they can't all get it right. If they're all giving us something different, somebody may be right, the others are wrong. Even when classifiers accurately classify the environment, they still may misinterpret the user's listening needs. So, for example, users in a social setting, the hearing aid might prioritize speech enhancement under the assumption that the conversation is a primary goal.

But maybe they prefer to listen to the background music. And we'll have that example in a moment. Okay. So it can't necessarily always anticipate their need, their needs, or their intents. Speaking of needs, again, one of the biggest challenges is that even though that the hearing aids may collect correctly classify the acoustic environment, they don't inherently know what the user wants to hear.

So this is the distinction, this distinction between the classification and listening to this is what the important thing for improving listening in real world situations. Let's consider two individuals in the same noisy restaurant. Both have identical hearing loss, both are using the same hearing aids. However, their listening needs are entirely different. So for example, this diner here seated at a table, he's primarily interested in focusing on a conversation with his companion.

It he's likely to benefit from strong noise reduction, suppress background noise and help improve the overall clarity of the speech. So the idea is get into more of a directional mode, minimize these distractions. Now contrast this with the waiter who's navigating the same environment, same hearing loss, same hearing aids, he has different needs. He's going to need to maintain an awareness of multiple sound sources simultaneously. He must maybe have to listen for a customer request, monitor conversations from different directions, remain alert so he doesn't bump into other people.

In this case, A more omnidirectional or an omnidirectional microphone setting was probably going to be better to help create that situational awareness.

So effective directionality then requires a careful balance, as we see in this teeter totter, between two competing needs. First is directional focus on the left, which prioritizes sounds coming from a specific direction, such as a conversation partner. The second is environmental awareness, which again, as we said, helps maintain a broad sense of their surroundings. Situational awareness. So if we go back to our restaurant vignette, the diner here is going to benefit from some degree of directionality when communicating with his dinner companion.

Directionality may take the form of traditional directional processing, which is often referred to as monaural beam forming, which is with the two microphones or binaural beam forming, which the two hearing aid microphones interact then to create a more narrow beam of directionality. Now, if we contrast this with the waiter, again, as we identified, he's going to have different needs. In order to do his job effectively, he must maintain a broader perspective, forming what we call a mental spatial map of his environment. This allows him to remain aware of the different sound sources and maintain some sense of psychological comfort just being have this awareness of sounds coming from around the surround sound. The ability then to create this spatial map is going to depend upon preserving or recreating monaural cues in addition to the pinna from coming from the pinnacoustics, as well as binaural cues, which we know low frequency interoral time differences and high frequency interal level differences to help then create this awareness for sounds coming from different directions.

These are the binaural cues as well as then considering different. Different form factors and how the hearing aids are then ultimately coupled to the ear canal. All of these are important to consider when designing these systems for effective directional focus in contrast with environmental awareness. So now let's consider this scene. It's a busy Nashville bar, lots of sound sources.

This is a unique challenge because there's. Now we have two individuals here, okay? In this case, we have these individuals, okay, who are trying to have a conversation. Their primary goals, just as before, is to hear and understand each other despite the competing background noise. For them, the ideal processing is going to be to suppress the background sounds, do some speech enhancement, minimize the background music.

Now let's take this other individual who's sitting nearby. Same hearing loss, same hearing aids. Now he's going, he doesn't want to hear those guys over there. He wants to focus on the music classifier. Again, might arrive at the same solution, but two different intents in how they want to listen in that environment.

That's the challenge is how can we adjust the hearing aid? How can we account for these different intents in these identical listening environments? Environments.

That's where there's idea of adding additional sensors to the hearing aids come in. So microphones are one sensor. Okay. They're bringing in acoustic information. However, we're starting to see more integration with motion sensors to allow the hearing aids to adapt more effectively in real world listening environments to both acoustic and motion based data.

And so let's just. So if we just look at then motion sensors. So motion sensors, what can they offer? So they can detect user movement, provide additional context for the information coming in from the microphones. This could be accelerometers that measure speed and movement that would allow the hearing aid to recognize if the user is walking, running or if they're stationary.

Gyroscopes can be used to detect head turns, which can help the hearing aid then to dynamically adjust the directionality based on how the user is engaging in that environment. Then we have magnetometers. We don't really, I don't know if these are actually implemented. I've seen these mentioned in the literature. These are like, if you think about like digital compasses.

Whole idea though those by combining acoustic and motion data, the hearing aid can then better make adjustments in real time to respond more effectively in a multitude of environments for different individual listening needs or intents. For example, if a user is on a noisy street, the hearing aid may apply a different type of noise reduction than it would if the user were seated at, at a quiet cafe. They detect that they're turning their head a lot in conversations can automatically adjust its beam forming to either broaden the scope, open up the directionality or narrow it down accordingly. The other thing too is that with all of these, again going back into the rd, that's a lot of information. So how do you train these things up so that they're most effective?

That's a lot of data. Right. So the idea here is again using AI behind the scenes in the research and development to integrate the motion based data with the acoustic information to make these inferences that we'll be talking about.

So again, this is like the, in this case, with this example here we have accelerometer measures. Acceleration of acceleration forces in three orthogonal planes can provide a wealth of information about whether a user is stationary, running or cycling. So here we have different patterns of motion across different activities. So in this example, the Patient is sitting still, the acceleration data is going to only show small variations in movement over time. Persons who's running is going to show large periodic fluctuations.

Acceleration, and if they're biking, may show moderate fluctuations with much less periodicity. And all of these distinctions then can help understand how the patient is engaging in that particular environment. So let's just consider the situation, the restaurant situation again. So again, we have all of that. In this case, we have the patient, the diner who's sitting down.

The hearing aid will detect that he's sitting, he's at a static posture, and maintain a strong focus on the conversational partner. If he gets up to pay the bill, the hearing aid will detect his movement and shift to a more omnidirectional mode to accommodate those changes in his listening needs, just by virtue of standing up and moving. And then we got the waiter, of course, just by virtue of the fact that he's moving around, detecting that, sensing, hey, he's going to need to have this more surround sound in order to be able to navigate the environment in the multitude of sound sources. Again, so using that information, let's take another common example. Walking with a conversational parking.

Most people naturally walk side by side when talking, talking and walking. Okay? So in this case, just by. Again, the optimal solution is to open up, go into a more omnidirectional mode rather than a direction of mode, so that again, the patient here in this case can hear his daughter to the right, for example, while they're taking a walk.

How about intent? You're hearing a lot about intent lately. That's kind of an abstract word. So I think some of the previous examples are a little bit more concrete. We can envision how.

That's kind of some very basic examples of a person getting up and moving in their environment and therefore having a different demand, different listening needs. What about intent? Because that gets something a little bit more psychological. And so the idea here then is to look at, at head turns and to be able to analyze head turns and try to infer what a user is trying to listen to. So let's consider a.

Here we go. Another very common. Like to use these dining examples. Okay. We have a diner at a noisy restaurant, and if turns his head to hear conversational partner, okay, this action, just by strongly turning your head, maybe it would indicate that he wants to prioritize speech from that, that direction and reduce background noise.

If it's just traditional beam forming, okay. The idea there is, is that they may not be able to detect that there's A change in that movement. So as if this person's, for example, turning their head, scanning the environment. Well, maybe having a narrow beam of focus is not going to be ideal. They're going to want to open it up.

Up. Let's take a look at this situation. So now in this case, now this individual here, he's trying to. He's the individual with hearing loss and. And very, again, noisy background.

But by virtue of the fact that he's looking at the music, in this case, the music is the thing that he wants to listen to and to then be able to again, maybe have a much more narrow beam of focus. Trying to infer that that is the thing he wants to listen to and not the speech that's going on around him.

Oh, this is a good one. Let's think again through again. I can mention this with frequent head turns, he think of like a Thanksgiving, big family Thanksgiving meal here. They're big fam. People are seeing each other, haven't seen each other in a while.

Okay. And the idea there is, obviously there are some situations for trying to hear the person across from you. Having that narrow beam of focus is going to be bad. So if you kind of lock in on them to try to reduce the side chatter. Okay, but then what if the conversation shifts, okay.

And all of a sudden you're trying to. You're doing a little bit of trying to figure out, okay, who said what. Okay. And sensing that this idea again, the patient's scanning their environment. Now it's time to open it up so that they then can take in the sound sources coming from different directions.

They find what they want to look at, look or listen to, and then stop scanning and then maybe narrow the beam up. Focus.

How about this one? So subtle, subtle head movements too. So this idea is. Well, just by using that same information, the same sensors can infer listener intent. So by just a subtle tilt of the head or if the user leans in, that sometimes indicates a increased listening effort or desire to hear conversations.

So in that case, the proper action would. May be to then have a more narrow beam of focus, much more aggressive noise reduction, maybe a little bit extra amplification, because that indicates that person is struggling in that environment. They're leaning in, trying to hear people better. Okay. So the idea is taking in this additional sensor information, incorporating it with the acoustic information to make more informed decisions about what it is the patient's trying to listen to in that particular environment.

Okay. Now, as we go through the cycle here, we're kind of at the point where we already kind of highlighted a little bit about user interventions ideas. Even with all of this sophistication, the solution arrived at the hearing aid, whether it's traditional processing, whether it's AI processing, may not quite accurately capture what the patient wants to get out of the environment in terms of their needs or their intents. So the idea here is can we use AI to help better interpret the user's interactions with the hearing aids? So this could be a couple of different forms.

So it could learn from more of a behind the scenes sort of approach. So whenever the user increases their volume modifies, noise reduction changes the directionality kinds of. We indicated with this team of audiologists or this audiologist following the patient are on all the time taking note that the patient is making these adjustments in these particular environments and to learn those patterns. Okay, that hey, the solution arrived is not the correct one. Let's see if we can maybe predict it better or arrive at a different solution based on what that patient wants.

And again, the idea here is that over time, ideally, that is the patient keeps interacting with their hearing aid, that the hearing aid will be a better may be able to make the better predictions for how that patient's going to want to alter their settings in a new environment. Maybe not the same environment, but maybe based on a new environment it hasn't seen before. So the extent that it can learn from that would be beneficial. And then we have potential solutions where the user is actually explicitly working with the AI to find an optimization optimal solutions. So the one I have in mind is this interface on an app gives the patient different comparisons, different settings in a B comparison format.

Listen to this. This. Which do you prefer? How much more do you prefer? It based on that selection will give the patient two different settings.

And it's basically trying to find an optimal solution for that patient for that particular learning environment and allows them to basically kind of hone it down, gives them the ability then to save that unique hearing aid. Setting this idea there then again is being able to learn from their interactions. And then to the extent that these hearing aids then connect to the cloud, in this particular case, the user can consent to having their information sent to up to the manufacturer for further analysis so that it can arrive at better solutions for the next person when they're in that situation based on how that person had preferred their settings for that hearing loss in that particular environment. That's this idea of crowdsourcing that I mentioned earlier. Taking large data sets from a multitude of users and being able to consider things like again Degree of hearing loss, style hearing aid, the coupling to the ear, the environment, okay, so on and so forth, taking all these bits of information.

So I could have a much more refined solution for that person next time for. For the. That another person in the similar situation or a new situation.

So it's very similar before. So this case, now the user is engaging, and this time it's more behind the scenes. The user engages. In this example, they engage the AI through the app or tap on a double tap on the hearing aid, and then the hearing aid kind of does its thing. So the AI then is basically refining its parameters, for example, volume, noise reduction, speech enhancement, directionality for that particular environment.

So the patients not necessarily have to do anything. They just have to initiate that, that sequence that then happens behind the scenes. It's not as iterative as the previous example. It's much more passive.

Then that brings us to analysis again. So now we have all this data. We've gone through the cycle, implemented a. We've designed some algorithms using AI with large data sets. Okay, we either we hard code our solution onto the hearing aid, put it in the automatics, or maybe there's some learning, some AI actually happening on the device.

We give the patient an opportunity to interact with that, to give it some feedback, to give it some output, some feedback to its output. So ideally then it can learn over time. And then the question is, where is that data analysis coming from? So it could be then coming. So be in the form of improving the human performance of that hearing aid, for example, being able to arrive at better recommendations next time for that patient or for the next patient, or again, having all of that information fed to research and development and going back into the R and D so that maybe in terms of future generations of products, understanding what some of the limitations are and the need for future improvements to keep that iterative process going.

And in terms of. It was starting to see some other applications mentioned in terms of things that are happening behind the scenes. So AI is seen one manufacturer, in terms of talking about using AI to optimize ear mold and dome selection to be able to have better recommendations for user preferences based on the degree of hearing loss, looking at the feedback there is in terms of what was recommended to the patient and what the patient was ultimately fitted with. We've seen some applications in terms of using AI to prevent performance degradation. So by continuously monitoring things like streaming stability, battery life, okay, Being able to detect early signs of instability or where the idea is that maybe the manufacturers can, as I discuss, correct these problems before they become noticeable, for example in the form of updates.

So likewise, doing some proactive maintenance by identifying trends across multiple users, for example, the specific model and the battery showing premature signs of failure on a large number of devices would then better help them to proactively implement a solution for that before the problem becomes much more widespread. And ultimately the idea then is to extend longevity and reliability of the hearing aids using the AI that's available.

Hearing aids are no longer being used to just improve hearing. We've seen different applications in terms of some health monitoring applications, in terms of tracking health tracking steps, tracking physical activity, social engagement. The idea is taking a much more holistic approach. Fall detection is a big one. The idea there is that the device head worn accelerometers may be as good or maybe better than the traditional neck warrant devices in terms of being able to detect if a person's falling beyond physical monitoring, they can become like a personal hearing coach.

We're seeing some applications here with the smartphone apps acting like digital trainers, tracking the usage, encouraging engagement and providing audiologists with more data driven insights in terms of how the patient is using the hearing aid, how often, what sort of environments, how many times they're engaging in social activity. The idea is again the whole idea is to have much more holistic approach to patient care, whether it be just strictly through the app or engagement with the audiologist. To the extent that hearing aids are integrating with smartphones, the idea is allow them to give this sense that the hearing aids themselves are becoming smarter. And just by connecting with all of the AI that's available on the smartphone, this could the voice command, cloud connectivity for other sorts of features like language translation. That's pretty exciting to be able to real time translate of course happening from the hearing aid to the phone to the cloud and back again.

But gives the sense then of having the hearing aids themselves being smarter but by application to the connection with the smartphone point and then finally this kind of I think the hopefully maybe the most important information where what does this all mean for you? That's there's a lot that's being discussed. Okay, I've discussed a lot of use cases, a lot of trying to motivate the need for AI in different applications in terms of listening intent, listening needs, trying to balance, you know, these are fairly generic situations are recognized. Okay, but how are we then to take all of this information and make sense of it? So the idea is there's at least there's some signals still some significant challenges that need to be addressed.

So one of the primary challenges is still I discussed using head turns and movement for user intent. Well, there's always going to be that limitation. I mean they don't always know what a user wants to hear, even with the more sophisticated sensors, as we discussed. And the concern there is if they misinterpret the situation. You know, we know the cost of misclassifications can sometimes be pretty significant if it applies the wrong solution and may cut out the patient from the things they really want to hear.

Maybe the, the speech for example. And so they get the idea there is, there's still room for improvement. Another is the ability to rapidly change, to adapt to rapidly changing environments. As far as I know this still I've read a anecdote from someone listening to, using some of these hearing aids. I think it was actually at the.

The can't remember the anecdote was, is that they initially were cut off from the, from the conversation first upon first usage because there was so much going on at once. And then over time it got better at being able to amplify the person they were conversing with. And so to a certain extent there's always going to probably be some sort of catch up up when the environment rapidly changes.

Data privacy concerns. We've heard of a lot of breaches over the years in terms of people's data, health data, financial data. Okay. And if patients are being told that their information is, or they guess, they consent. But how is my information gonna be used?

I mean, you know, the idea here is we're using these hearing aids to have a much more personalized solution that's giving a lot of information about that person perhaps. So how is that information going to be preserved? How is it going to be privacy be maintained? So just point is be aware of this. Patients will bring this up.

Energy consumption, battery life. Manufacturers are still balancing this. Okay. The idea here is even with rechargeability, there's going to be an overhead limit in terms of how much processing can we do above and beyond amplification, above and beyond streaming. If we're going to be streaming and amplifying and doing AI all day long, how much battery life is that going to preserve?

Is a patient going to have to recharge their hearing aids on a much more frequent basis? Challenges that here is this data quality, this data in particular is on the training sets. Okay. So how do you get this multitude of data sets, these tens of thousands or millions. Okay.

And try to maintain, be aware that there could be bias in your data. And for example, not having A wide variety of environments or talkers with different accents, for example, or representation of age. I know, for example, the voice assistants at home, they struggle with my kids voices for some reason and it's not because they have articulation problems, it's just it makes much more frequent errors. And I think it's probably because those things were not trained on children's voices. So having good quality data at the input in the R and D site.

All right, so as these continue to evolve, as AI is continuing to evolve in hearing, it's important question, are these advancements truly making a difference or are we still waiting for real world validation? While it has the potential to transform our everyday fitting experience for patients, there's still gaps in research and evidence that needs to be addressed. One of the biggest ones is the limited amount of independent research. These devices set at the outset are often marketed as improving speech clarity, reducing effort. But all of the data, most or all of the data comes from manufactured LED research.

So while these studies, they are somewhat insightful, we always have to be concerned about the bias. Do they necessarily reflect real world listening conditions? So what we need are more large scale independent studies to really evaluate the impact of AI in everyday situations, above and beyond what's advertised by the manufacturers. Related concern is the real world effectiveness is unproven. While the

AI and hearing aids demonstrates impressive performance in controlled lab studies, their success and dynamic unpredictable situations is not as certain.

So if we have multiple conversations, background noise, individual listening preferences make the real world hearing much more complex than what might be simulated in a controlled laboratory situations. And to the extent that AI can't account for all of these variations may then lead to underperformance in everyday use.

And so then we have some differences. So what are you, you're going to be seeing or seeing lots of advertisements. What are some ways that AI may differ between manufacturers? One is this dedicated chip AI chip. So they have two chips in the hearing aids, one's for all of the regular signal processing, one for the AI or is it integrated?

There's another, at least one manufacturer integrates the AI inside the chip with the regular processing. And the big difference there is in terms of the complexities they're able to handle. So the dedicated chip being able to handle maybe more complexities than the integrated approach, but the integrated approach they're advertising as being able to have much more, it's much more computationally efficient, it's much more, it's better for preservation of battery life. The idea is that you can run AI all day on A single charge, not have to worry about recharging your hearing aids. Whereas another manufacturer has a dedicated AI chip.

And I believe they limit you to three hours of the AI once you configure, once you factor in, patient may want to stream for a certain number of hours. And of course you have everything else the hearing aid is doing in order to make sure that the patient can get through a day without having to recharge their hearing aids. And that's this idea of power consumption, of battery efficiency. Those two could vary independent of whether or not it's an 8 dedicated or integrated chip. But those are two differences.

Another difference is the model that's used. Again, we started this off machine learning, deep neural network learning. Again, the idea is machine learning is going to be much more computationally efficient, but it's going to be limited in terms of its ability to adapt and learn over time. Talked about the data a little bit throughout this. This is the training data, how many samples, the quality of the data, the diversity of speech background samples, the authenticity of the data.

Is it just lab recordings through real world samples? How much involvement was there in terms of some human oversight to make sure that the labeling of these environments was accurate? And as we discussed, biases in terms of age, gender, ethnicity, being able so that again, whoever's using these, that these models will be able to account for these variety of situations. So these could be some differences in terms of the quality, if you will, of the AI from different manufacturers. Motion sensors we talked about, some are integrating motion sensors, others are not.

This ability to interact with the outputs, the hearing aid makes a decision, the user gives some feedback to integrate their preferences. That's a difference you will see between the different manufacturers. Again, where is the processing being performed? Is it being performed in the cloud for later use? Is it being performed on the smartphone?

Maybe a little delay. And how is that integrating with the actual processing itself? Is it just offloading solutions onto the hearing aids? Where's the device itself actually doing some sort of AI on the chip or an integrated chip? These are all differences that you will see.

And then finally, what do we do with all this? So our role is to help patients understand and navigate this new world that we find ourselves in. And they're going to be aware of AI probably being exposed to these advertisements. So we kind of have to balance enthusiasm with realistic expectations. So we've been here before, you know, every new technology and hearing aids is revolutionary and it's going to solve all the world's problems.

AI is often marketed in the Same way promising dramatic improvements. However, again, independent research is still lacking to validate many of these claims in real world environments. So as clinicians, we must encourage critical thinking, advise our patients not to accept AI marketing at its face value. To the extent that we can, it's still limited. Rely on peer reviewed research when and if it ever comes out.

Okay. Manage patient expectations. So patients need to understand that while AI can assist their hearing aids, it's not going to replace the fundamentals of amplification, directionality, noise management. So it's not going to be able to completely eliminate their difficulties. So the point here is to set realistic expectations so they can appreciate the benefits without being frustrated when it when they do and eventually experience some of its limitations.

Those limitations are going to often be again, complex noisy listening environments with multiple talkers, overlapping conversations. We should not be surprised if AI still struggles to isolate, prioritize the speech of interest.

We must then make sure that we have a patient's best interest in mind, so remain in control. I understand that we must remain at the forefront of patient care and guide the integration in a way that enhances, not replaces our professional decision making. AI is a tool, it's designed in its appropriate place, can be used to support our everyday clinical decision making. Never replace basis. Allow for much more personalized fittings by taking in that information that we can gather from data logging, for example.

Again, be aware that patients might come to the appointment concerned about privacy considerations. Where is my information going? Is it going anywhere? Is anybody selling my information? A lot of that's tightly regulated, particularly in the eu.

So to the extent that these manufacturers are being manufactured and sold there, a lot of those considerations, if they abide by those laws, should be in place. So be aware of how to address those concerns. And of course patients, just like with any technology, it's not going to be for everybody. They're going to have different levels of tech savviness. And so again, this is no difference.

AI is not AI. And the hearing aids may not be for everybody, particularly if it involves user interaction to help train it up and fine tune it, stay up to date, as always with everything in the field, nothing's different there. Again, offer AI is an option, not a necessity, nothing new here. We always teach our students, you validate your outcomes. This is no different.

So we should, if we're just doing best practices, a lot of this should come natural. A lot of this should be new information in terms of going to, examining the weight of the evidence and then again evaluating, you know, the long term Effectiveness using our tried and true validation measures to actually have some data and whether or not these things are making a benefit for our patients. Just some analogies you can use for different types of patients, some different metaphors and then the last. Okay, just as we kind of already been talking about in the last couple of minutes. Okay.

The idea here is that while AI can provide valuable automation, important point is that it supports, it doesn't replace our clinical expertise. Okay. Patient care. It's always it's crucial at what we do. So you know, we are essential in guiding the use of AI, interpreting the results, ensuring the best outcomes for our patients based on their unique needs.

AI doesn't have the ability to replace that to take into consideration their preferences as even though that we've discussed to a certain extent accommodating preferences intents. But we still need our human expertise. And ideally then is this should be seen as a collaborative partner, not an independent solution. We use it to complement our clinical knowledge, our professional expertise and to main be in control of the process. And so point is, you know, as good as it's going to get, it's going to lack the ability to do what we do best which is address the emotional, social, psychological aspects of their hearing loss and overall help rehabilitation their hearing needs.

So thank you so much Dr. Alexander for your time and expertise and thanks to everyone who participated in today's course. We look forward to your feedback on the course evaluations and hope to see you in future courses. Please note that the exam will be available in your pending courses on our website. Fifteen minutes following the completion of the course, you will have two attempts to earn a passing score and must take the exam within the next seven days.

This concludes today's course. Have a great rest of the day.