Explaining True Focus and True Noise Reduction

Presenter: Laura Schachtel, AuD, FAAA

Hi everyone. Welcome back. Hopefully again, I'm Dr. Laura Schachtel, the product training development manager for Beltone. And today I get to talk about a new noise reduction feature that we call true noise reduction.

And this is the first time in my tenure with GN, so probably about 15 or so years that we've actually introduced any type of new noise reduction feature. And if that wasn't enough so new noise reduction, we. We also introduced a new beamformer with our Beltone InVision 17 products. So I'm sure after hearing that Beltone is now using DNN in their hearing aids, you all quickly went to your computer and typed in what is a dnn? And were able to easily figure it all out.

And kind of surprised you're all joining me today because I think when you do a quick search, it explains it pretty easily. Being very sarcastic because what you probably got is something that looks like this picture on my title slide here. And what I don't understand what this is doing and how does this relate to what my brain's actually doing as well. So I will hopefully kind of demystify that and make some sense as to what a DNN is, how we're utilizing it to help with hearing in noise for our patients. As Kimberly stated, there's lots of handouts, so please feel free to download those.

If you downloaded them from yesterday, then you're good because they're the same handouts that will be up there all week. Except there might be a different note handout for this presentation. Again, beltoneuniversity.com also has all of these handouts as well. So let's get on to our learning outcomes for the day. After this training, hopefully you'll be able to describe how Beltone's Intelligence Augmented approach to Artificial Intelligence empowers patients to take control of their hearing experience.

Be able to compare True Focus to Clear Focus to Ultra Focus 2 and Ultra Focus. Be able to explain the differences between True Noise Reduction and Sound Cleaner Pro. And then lastly, explain how Beltone achieved a balance between DNN chip efficiency and battery life. Now, these are the risks and limitations of this course, so please go ahead and review those as needed. But let's conn to dive in

now.

So before we really dive into AI and DNN technology, I want to explain our approach that we call Intelligence Augmented. Now, we understand that there is nothing more powerful than the human brain, and we want to empower our patients by providing them with greater control over their hearing experience. Intelligence Augmented is all about empowering the human over the machine. To deliver the best outcomes. So imagine a beautiful day where a grandmother and a granddaughter are riding bikes together, sharing laughter and joy as they create some cherished memories.

The sun shines brightly and the air is filled with sounds of their giggles. However, as they ride along, they encounter a steep hill. The granddaughter, full of energy and excitement, really wants to conquer that hill. But the grandmother feels kind of a wave of fatigue wash over her. She knows that riding the bike is going to be much easier for her than walking the bike up the hill, but she worries about her ability to pedal up that incline.

In this scenario, we can liken it to an experience of many of our hearing aid users who find themselves in noisy environments. While hearing aids can help them engage in conversations, it requires significant effort and can be quite draining. Now, let's introduce a twist here. What if the grandmother's bike was equipped with an electric motor? This addition would empower her with the choice to activate the motor or keep it off, depending on the terrain's difficulty.

When she decides to turn on the electric motor, she is still maintaining complete control over the bike. She can pedal, brake, and steer as she wishes. As they approach that daunting hill, the electric motor provides the necessary boost to make the ascent much more manageable. This way, the grandmother can focus on enjoying the ride and the precious moments with her granddaughter Rather than struggling against the physical challenge of the hill. The motor acts as a support ally, enhancing her experience without taking away her autonomy.

In this analogy, the electric motor represents the additional support that hearing aids and deep neural network and AI can provide in challenging situations. Just as the motor helped the grandmother navigate the hill with ease, hearing aids can assist users in engaging in conversation amongst the background noise, allowing them to participate fully in social interactions without feeling overwhelmed. Ultimately, the grandmother and our hearing aid users can enjoy their experiences more fully empowered by the right tools to help them along the way. Again, that bike analogy kind of showcased really what all of our patients are going through. So many of our patients have a sensorineural hearing loss.

Again, with the sensorineural hearing loss, the hair cells in the cochlea have been damaged or destroyed due to aging, trauma, noise exposure, or ototoxic medications. This results in a reduction in both the quantity and the quality of the neural signal that the inner hair cells send to the brain. As hearing loss progresses over time, that neural code continues to deteriorate. And as research has shown us, prolonged auditory deprivation can start to lead to brain atrophy. And that's not a good thing.

Now, with AI and DNN Technology, along with all the other features in our hearing aids, we can help our patients hear better. But these features and new technologies still cannot fully adjust to a person's listening preference. With our approach, we've combined the brain's natural processing power with AI technology. When the extra help is needed, just like adding an electric motor to a bike with a hearing loss, our patient's brain must exert more effort to listen and allocate additional energy to decode and understand the compromised neural signals. Utilizing AI and DNN assists the brain in processing incoming sounds more efficiently.

So in simpler terms, AI enhances the clarity of the signals for the brain to interpret, thereby reducing cognitive load. Again, our goal is to help the human, not replace the human. So to state this a little bit differently, we know there is a lot of noise out there, which is why hearing and noise is our number one

challenge that our patients face. To ensure our neural core processor, sorry, our natural core processor. Our human brain gets the help it needs when it needs it.

We introduce two co processors that are embedded in our new dual chip. Using our 360 chip, we begin to classify sounds and start to filter out some of that background noise. With our new DNN chip, we can create greater contrast between noise and speech, making it easier for our human brain to make that final decision on what to tune into. Now, before we get into our new noise reduction feature, I just want to quickly review the updated beam formers that each of our Beltone Envision technology level has when you're in that hear in noise program. So with our 17 level micro RIE you have True Focus and that uses our new DNN noise reduction that we call True noise reduction.

And with the 63 or 64 style you have Clear Focus that uses our legacy Clear Clear Sound Cleaner Pro. And with a 9 level with any RIE style you have Ultra Focus 2 using Sound Cleaner Pro and with an Envision 6 level again any of the RIE styles you have Ultra Focus again using Sound Cleaner Pro. And then in our envision 4 level you have Smart Beam and that includes the beam steering where depending on the signal to noise ratio, that beam width will go from wide to to medium as needed. And again that is still using Sound Cleaner Pro. Again with our new 4 mic beamformer, in both Clear Focus and True Focus we change the frequencies that are included in the beam.

So by removing that low frequency omnidirectional response, less noise enters the hearing aids from the sides and behind which which reduces the overall noise level and helps maximize the efficiency of our noise reduction feature. So if you remember with Ultra Focus 2 and we kind of looked at the same graph, we had the lower frequencies being processed in an omni fashion. Those middle frequencies were our 4 mic beamformer. And then on this slide, just as you see it, those higher frequencies were being processed in a fixed directional format. So what we did is we said let's put our beam former, or let's allow more frequencies of our beam former, more frequencies to be involved in our beam former.

And so we move that down into those lower frequencies. I spoke about this yesterday, but research has told us that when we're in those moderate to loud background noisy conditions, having that omnidirectional low frequency response really doesn't add much more benefit. And the benefit really is in having more frequencies in that directional beam. And so this is the upgrade that we've now created. So let's now move on and talk about AI and DNN and understand how we're using this technology to help with our noise reduction.

So I'm sure we've all heard about AI and whether that is something good or bad, we all have ideas about what it is. In the most general terms, AI is the ability of machines to simulate human behavior and make decisions like humans. But there are smaller subsets of AI that you may have heard about, like deep learning and machine learning. So with our true noise reduction, we are using deep learning, where we trained an algorithm to make decisions. And we've used this kind of deep learning in the past when we trained our HearMax app to recognize ears, when we implemented our check my fit feature.

Now the next step in this progression is something called machine learning, where the machine, or in our case, the hearing aids, would automatically make a programming change without being programmed or told to do so. Or you could think of it as the hearing aids are actively learning and adapting. Now, right now there is no hearing aid that is actively machine learning. I think we're all still in this deep learning phase. But could it be possible in the future that machine learning takes over?

Probably. And that could be the path that we are now starting on. So again, I assume when you heard about dnn, just like I did, you all typed in what is a deep neural network. And again you got something back that probably looks like this, which is so helpful when trying to understand this kind of random computer based algorithm thing. And the search engine will tell you that it looks just like this because it's mimicking the normal process of a human brain.

And that's really all the explanation and we can all go on our way. It's really not that easy. But what you can see that's easy is that there is an input layer and it gets filtered and then sent out in an output layer. So for us, we'd want to input speech in noise, have it filter through the system, and then at the output we would want just noise, sorry, just speech with a little bit of noise. We don't always want to eliminate all the noise because there is some importance to noise.

But put in speech, in noise, get out, essentially just speech. Now. Now the confusing part are those middle hidden layers or those blue nodes in this diagram and then all of those lines like, what does all of that mean? So my global colleague came up with a very rudimentary but pretty simple, easy way to explain what is happening in these hidden layers. And this might look familiar to some of you, but to explain what a D and N is, I'm going to talk about the game Guess who.

So I remember playing this when I was maybe like 7 or 8 years old. Now, if you've never played this game, it's pretty easy. It's a two person game. There's a deck of cards. Both players randomly pull a card and what basically the opponent is trying to guess who your mystery person is.

And then you have this board of all of the faces like you see in the background here. And you will kind of ask, does this person have a certain attribute? And if the answer is yes, you leave that tile up. And if it's no, you knock down all the people that do have that attribute. And I honestly never knew, of course, you know, seven years old and DNN really didn't exist in my vocabulary, but I was actually learning what a DNN was all those years ago.

So the first layer of notes here represent the personal attributes that distinguish all the players in the game from each other. So it'd be things like wearing a hat, has a mustache, wears glasses, has blue eyes or brown eyes, has long hair. And then that second layer of nodes, we have our possible suspects or our mystery people that we're trying to guess. And so now we start to play the game. So we'll start with our first attribute.

Does our mystery person wear a hat? And in this case, yes, that person wears a hat. And then we would map that all to all of our possible mystery people. And the answer then gets coded because it's an algorithm, so it's all math, gets coded with a probability function between 0 and 1. And so a 0 means that there is 0% chance that that person has a hat.

And you can see that the second gentleman down with the orange hair clearly does not have a hat. And that's why his line is black, because 0% he has a hat. But all of those blue lines get a number one coded to them because we are 100% sure that those four, five people do have a hat. And so they get that number one. And again, this is just an algorithm.

So it doesn't actually see a hat. It doesn't actually see the face. It just sees a one or a zero. And so we're training it to accurately put that one or zero for that physical attribute. So then we continue to play the game and we go through each attribute.

So does the person have a mustache? And we map to mustaches. Does the person have glasses? And then we map to the glasses. And you can see now that even though this is just two of those hidden layers, this is now starting to look like that diagram of what a DNN is when you search for it.

So all of those connection lines are basically mapping an attribute to the possible outco that we want in the output. And so we'll keep going with blue eyes, and then we'll map out the answer or code that answer into our dnn. What about brown eyes? And we code that to our dnn. And then lastly, hair length.

And so we go on and on until we kind of get to the end here. And then if you've been playing along, hopefully you can pick out who our mystery person is. Our dnn, as it's kind of playing this game, it's simultaneously able to extract all of this information very quickly and so that it can call out that our mystery person. And again, maybe you found this as well. This is the person that has the most blue

lines drawn to them is Claire, because she has the strongest connection to all of the attributes, as she has three blue lines and everyone else has one or two.

And again, this example is very basic, very rudimentary, but this really is what the DNN in our hearing aids are doing. There's just more layers to it. And so the more that we play this game or the more that we train our dnn, it uses its previous experiences to make better decisions and to come to a conclusion even faster. And I just want to assure you that once we've fully trained our DNN chip and we put it into our hearing aids, no more learning is taking place. It's just a chip like our 360 chip doing the job we have told it to do.

It is no longer learning. It is not getting smarter. It is not getting better. It is what it is when we put it in the hearing aid. So if you have any patients that are worried about AI in their hearing aids, that the hearing aids will be spying on them, that whatever their worry is, you can just let them know, Nope.

It's just like any other algorithm, any other feature in our hearing aid. It's just been trained a little bit differently. So here is an example of what is happening in the DNN chip, actually in our hearing aids. So the input corresponds to the energy in each of our hearing aids. 17 bands.

And this could be a mixture of speech and background noise. And the DNN then looks at different parameters, kind of like those physical attributes of our mystery people. So it's looking at instead of hat and mustache and eye color, frequency modulation, amplitude, spectral balance, waveform shape, time variation. And it tries to extract those sounds and understand how many sounds there are. And we've never before been able to extract so much information from a single sound source almost instantaneously across so many different parameters.

Now, the output of The DNN is 17 outputs, one for each of our 17 processing bands. And you can think of those middle hidden layers with all of those nodes and all of those lines kind of as there's like an

attenuator adjusting dial there, and it's constantly, by the millisecond, adjusting the contrast to figure out how much gain reduction is needed in order to maximize speech and reduce the unwanted background noise. I stated earlier when I misspoke, we don't want to fully eliminate noise because we know noise is important. It helps with maintaining alertness and just keeping a patient aware of what's going on in their environment. But we definitely want to try to reduce that down as much as we possibly can, because we know that speech is really the important thing in that environment.

So I played this video yesterday, but here it is again. This is just an overview of how our dual chip works. So the advanced technology enters the hearing aids and gets separated into the 17 distinct frequency bands, which again allows for more precise filtering of those unwanted sounds. Now, true noise reduction is operating on our DNN chip, while all other hearing aid features continue to be processed on our 360 chip. And so the output from each chip come together, they get blended together.

So those 17 channels blend right back together. And what we get in that output is reduced noise with enhanced speech understanding. And this integration ensures that our patients receive the best possible auditory experience. Again, reduced noise and focused speech and understanding. Now, you Might be wondering how our new true noise reduction feature differs from what Sound Cleaner Pro offers.

And Sound Cleaner Pro utilizes a technique known as spectral subtraction. So with Sound Cleaner Pro, the hearing aids estimate whether the incoming sounds are speech or noise based on their acoustic properties. However, it's important to note that this method doesn't truly understand the nature of the sound. Rather, it perceives an outline of the sound without capturing its full complexities. So again, it's just kind of getting a rough estimate.

It's not spending the time to really deep dive and really parse out is this noise or is this speech? Because there are some speech sounds that look like noise. So again, I've always trained that the hearing aids are listening. It gets kind of a snapshot of the noise in the room. With Sound Cleaner Pro, it then takes that snapshot and says anything that kind of looks like this is clearly noise and will reduce it.

And so we have the speech a little bit louder and the noise has been reduced down. So in practical terms, what Sound Cleaner Pro is doing is it's reducing that short term noise spectrum, kind of that snapshot from that total auditory signal. And the result is an output that consists of speech, again combined with that reduced background noise. Now, these reductions are applied based on an estimate of the signal to noise ratio at any given moment. So essentially, if the signal to noise ratio is poor, then we would have more reductions being applied to specific frequencies, similar to what this graph is.

And I'm not saying this graph is what Sound Cleaner Pro is doing, but it's giving you a kind of idea of what could be happening. And if the signal to noise ratio is better, then we're going to have less noise reduction. And while Sound Cleaner Pro has been efficient in many situations, which is why we still have this feature and we've had it for many years, because it works great. It's still applying generalized reductions based on a sound outline. It's not going in more strategically or precisely and cutting through all of that background noise.

Now, this approach can sometimes negatively impact speech audibility or sound quality, particularly in complex listening environments where clarity is critical. So if you attended one of our profit events last year, we played sound recordings of the different levels of Sound Cleaner Pro. And with the more aggressive levels, this feature did begin to negatively impact speech sounds. So the F, the V, the S, the Z, and the th sounds, those are all called fricatives. And fricatives are produced by having a turbulent

air flow through your mouth.

Turbulent airflow is kind of a noise source. So when we set Sound Cleaner Pro more aggressively, it is saying, aha, these speech sounds are noise. And it reduce down the speech sounds, which is really not what we want, because with I with normal hearing could tell that that signal was not as clear as I know it should be. You can only imagine what that's doing to your hearing aid patient who has a hearing loss, probably high frequency sinus reneural, and they're already not hearing these fricatives. And so now we're making it even more muddied for them because this feature isn't as precise as we would want it to be.

So to contrast that with true noise reduction, where we can actually see what that noise is and we know what that noise is, and again, we can more precisely go in and reduce that noise or peel away that noise from the speech. So when looking at this feature In Solus Max 2.1, you will see three levels of reduction. Moderate, the default, considerable and strong. And again, I've added in those maximum DB attenuation amounts for you. Again, there's two numbers.

The first is for those six lower frequency bands, processing bands. And then that second number is for our higher mid and higher frequency bands as well. And the larger amount of reduction that we can offer in those higher bands allows for better sound quality, meaning it should allow for those fricative sounds to be heard and not be muffled. And then you can see that based on these values for considerable and strong, we are offering more noise reduction than what you could get with Sound Cleaner Pro. So you might be questioning how we added a whole new chip into our hearing aids that does such a great job at reducing background sounds, but still have a hearing aid in the same micro RIE size and provides the same battery life as our previous rie.

Well, first we did this because we were very smart with how we trained our DNN chip. And our training process was very efficient. So we designed our chip using a recurrent neural network, so it doesn't

require as much battery power as maybe another way of designing a dnn. A recurrent network uses information from the past to make a more accurate and faster prediction of what's going to happen in the future. And this helps preserve battery life.

So kind of think of it as a baby learning to walk, right? They crawl, then they can stand, then they kind of take a couple steps holding onto something. Then they take their first tentative steps not holding on. And then before you know it, they're off running because they've learned through trial and errors. And eventually they figured out and got really good at it and they can just go off and run.

That is what the recurrent neural network is doing. Now remember these middle hidden layers with all those nodes we talked about? The more nodes there are in a dnn, the more battery power is used. So again, we were very smart with how many of those hidden node layers we used. Now, secondly, we trained our DNN with 13.5 million speech sentences in noise.

Not just words, but full sentences, with male and female voices in different languages, with varied vocal effort to get all the nuances of the speech signal. Now, the noise sounds were noises that everyone would agree are noise like traffic dishes, clanking, road sounds. Now, did you know that if you were to speak 13.5 million sentences, it would probably take you about 25 years? So our Beltone and Vision 17s have 25 years of speech conversations in them. And then on top of all of that, we are combining this with our newest Beam former that allows less noise in the hearing aids.

So the speech in front is the main focus. Now just kind of getting back, if we were to train to have trained our DNN with not sounds that everyone agreed to was noise, but more ambiguous sounds, say like waves crashing on the beach or the crackling of like a fire, then we might actually have spent more time kind of with wasting battery life and maybe over processing sounds. Because if you start to over process and kind of go, is this noise? Is this not noise? It reduces the gain for sound that the patient actually wants to hear or can start to impede audibility.

So by using those ambiguous noises, some of our competitors are maybe kind of stepping on their own foot in a way. So I just want to point out that it's really the quality of the sound samples, not the quantity, that makes for an efficient DNN noise reduction feature. Again, when we look at our battery, a typical use using our bell tone, using our new true noise reduction and true focus, about 20 hours of use, right? If someone goes in for a luncheon, into a restaurant, they use this for a few hours. Maybe they stream a phone call or two.

20 hours to me is still all day use because you have to sleep sometime. And again, if your patient forgot to move out of their hearing noise program back into their all around program, they still would get 10 hours of use just being 100% in that true focus with true noise reduction feature on their hearing aid. So just to summarize, with Intelligence Augmented, we are using three processors. So we have our 360 chip, the new DNN chip, and then the human brain again, as we want to empower our patients with the technology, not the other way around. Our new beamformer Clear Focus uses full band directionality which helps reduce background noises because it reduces kind of the sounds coming from the side and behind.

And you will see this directionality drop down in the Hear a Noise program when using a 63 or 64 style in our 17 level and you will see the directionality option True Focus and that will be using true noise reduction when using a 17 level micro RIE. True noise reduction uses a deep neural network to help remove noise from the speech. From its training, it learns from previous experiences to more efficiently and precisely remove noise as needed. Again, all from its training, it's not actively learning anything new once you put the hearing aids on your patient. And again by training our DNN with quality samples, not a larger number of samples, we have a smaller size hearing aid with longer battery life and sound that is clear and not over processed.

And so with that I just want to thank all of you for taking a quick 30 minutes out of your day to learn more about True noise reduction, a little bit more about Trufocus and really how we implemented this new AI technology of deep neural network into our hearing aids. I will stay on the line a little bit longer if any questions do roll in. Otherwise, if you have any questions you can always email me or you can take the corresponding Beltone University courses. They're on the homepage. We have one on Envision, one on Beltone Commence, one on the dnn, two on updates to the Hearmax app that I'll talk about on Friday and then one on the update to Solus Maps that I will also talk about on Friday.

So I'll let you guys type in, but if you have any kind of other questions or need some further training, please check out our Beltone University homepage. We have self paced modules on Beltone Envision beltone commence. There's two on updates to our HearMax app and then one on our update to Solus Max 2.1. So two questions kind of came in. I'll take a look at those.

Otherwise, if you do have to run. Thank you so much for joining me. I hope you all come back tomorrow to join my colleague Kathy Liston as she talks about our other new product, Family Belt Tone commence. So if you need to go, I appreciate you for joining me. Otherwise the two questions are, question is, will the DNN AI get better over time?

How would that work? I assume, I mean I'm guesstimating but I think what is happening or from what I've been told is all manufacturers right now who are using dnn, it's not their own developed chip. They bought it from someone. So just like our receivers, most of the hearing aid manufacturers buy receivers from one location. Everyone right now has bought a DNN chip.

So I think the next generation would be that each hearing aid manufacturer maybe, maybe develops their own DNN chip. Maybe they can develop it so that it's not two chips, it's one chips, and it works more efficiently. So again, reducing battery life. I think the next step off of that is to get into machine learning, where instead of saying this is all happening in our hearing noise program, we just say this is

happening. And the hearing aid is now toggling between the two programs that you provided your patients or the three programs that you provided for your patient.

So this is where I think it's going. I could be completely wrong, but that's kind of. When I look at my tea leaves this morning, that's kind of what it said. Another question.

Question. Would it almost be like remote care where it sends an update remotely to the hearing aid or is the DNN essentially fixed and it won't improve until. Right. So right now our DNN is fixed. It's just like any other feature in our hearing aid.

So let's take Smart Gain. If we were going to update something to Smart Gain, we'd have to then like send you a firmware update to then firmware update and then offer this update to your patient. So our DNN right now, in the envisions that I have sitting right here behind my monitor, it's a dumb dnn. It learned and it's not ever going to get any smarter unless maybe there's a firmware update, which I don't know about. But like with remote care.

Yes. If there was a firmware firmware update, then that could be something we could push out in that way. But for right now, what's trained on our dnn, it trained, it learned. It's burned in the chip and it's not getting any smarter unless we do something to then give it more intelligence.

Right. So possibly a firmware update. I have a feeling though, that just with the way a DNN works and the reason why it can only be on our micro, because of the rechargeability, it needs a very smooth, consistent battery. So I am wondering just now as I talk it through in my head again, I don't know anything. A firmware update remotely might be hard because of the distance and the WI fi and so I would think if we ever were to update DNN with the firmware update, it'd probably have to be one done in the office.

And I have a feeling there probably won't be firmware updates because it might be too hard to get that in. Or the package of the update might just be too large where it is just a new hearing aid. But again, I know nothing more than what I'm sharing with you. This is all just me kind of guesstimating and thinking ahead of what could be happening. Another question.

Oh my goodness.

So can the chip we currently have learn or get smarter? No. The environmental classifier in the Serene did a good job, but it was fixed in what it could do. Correct. So is this one.

So that's something that could be happening. Like so. Sorry, let me finish reading the question. So the environmental classifier, so smart gain and serene. It does a great job, but it is what it is.

So could AI with all of these inputs get more accurate over time and then kind of integrate all this together? Yes, that would be more of the machine learning, which again, we're not at machine learning just yet, so we're kind of dipping our toes into deep neural networks and as the one graphic show, kind of the next leap would be into machine learning. But you're thinking ahead in the right direction of what AI can open up for how we can really help our patients with our hearing aids. Thank you. Please, no more questions, but I hope I did answer them all again, if you have any more, please, I don't want to call you out, but please feel free to email me and I can try to answer any more questions that you have.

So I do see that we are a little bit over our time. Thank you everyone who stayed on to hear me more. Kind of predict the future of what our hearing aids might have if we could have firmware updates, what AI in the future looks like. Again, I am by no means the expert in this type of technology, but I have a feeling this is where we're going. So thank you everyone for joining me.

I hope you have a great rest of your day and I hope you all join Kathy tomorrow when she talks about Beltone Commence. Bye everyone. Thank you.