True Focus: How deep learning augments human intelligence

Jennifer Schumacher, AuD; Jennifer Groth, MA

Abstract

This paper discusses the integration of deep learning in hearing aid technology, specifically through a new noise reduction feature called True Focus introduced with Beltone Envision™. This system integrates an upgraded 4-microphone binaural beamformer with True Noise Reduction, a single-microphone noise reduction system which utilises a Deep Neural Network (DNN) to enhance auditory experiences. Unlike traditional single-microphone systems that rely on model-based approaches, this data-driven method mimics the interconnected structure of neurons in the human brain, allowing for adaptive and efficient sound processing. We advocate for an "Intelligence Augmented" approach, emphasizing the collaboration between human cognitive abilities and Artificial Intelligence (AI) to empower users, ultimately leading to more natural and effective hearing solutions.

Imagine you are a bicyclist standing at the bottom of a huge hill. You want to reach the top, and your bike can certainly help you get there, but it will be an incredibly exhausting ride. You might even give up depending on how difficult the hill is and how tired you may be feeling. Many hearing aid users might describe trying to hold a conversation in noise like a ride up that hill, their hearing aids being analogous to the bike. It's possible but can still be effortful and draining, yet trying to follow conversations in unfavourable listening conditions is where hearing aid users rely on their devices the most.

Now imagine your bike has an electric motor. It's your choice as the rider to turn the motor on or leave it off. When it is active, you still have full control of the pedals and brakes, and you're still the one deciding which direction the bike will go. But now that you need to ride up that big hill, the electric motor can make the ride much easier and take less effort, even with you in control of the other aspects of the ride. The motor augments your natural action of riding a bike rather than autonomously taking over control of the bike ride.

This example helps illustrate what makes our approach to hearing aid development different. We look for opportunities to combine what the person wearing the hearing aids and their brain do well, with what technology does well. While it is common to think of applying technology to help the user, an implication of helping is that it takes a level of control from the user. Our approach is to apply technology to empower the user. Combining the strengths of artificial technology with biological processes of the human is what creates our solution. We call this approach "Intelligence Augmented" rather than the more commonly used "Artificial Intelligence". Intelligence Augmented is an integral part of the overarching Organic Hearing™ philosophy that leads us to create products and solutions that sound natural, feel natural and allow users to connect naturally. This approach has formed the basis for our Organic Hearing design through the years and remains true, even as more advanced AI is integrated into our products.

Artificial Intelligence in our hearing aids should function in an analogous way to the electric bicycle motor. We want the technology to be a useful addition to solving the uphill struggle with hearing in noise, but it should not interfere with the hearing aid user's listening goals. The hearing aid user remains in charge of the listening situation. No matter how effective a technology may be in reducing noise, it will never be able to replace the processing abilities of the human brain. So many of the strategies we humans use to hear in noise go above and beyond just hearing – we must also divert cognitive resources to focus our attention on various sounds, decide what we are interested in, and use our brains to filter out the rest of the sounds as noise. Hearing aids help most with the first step in the process, which is improving audibility so the sounds can be heard.

Beyond audibility, it's up to the brain to make decisions on what to listen to. A hearing aid cannot replace these higher-level cognitive functions, even with AI. Instead, our Intelligence Augmented approach helps ensure the user is receiving the best possible representation of the sounds around them, so their brain can take that upgraded information and do what it does best – listen effectively.

Listening effectively in noise

It is well-known that hearing well in noise is a struggle for people with hearing loss and a main driver to seek help. Although hearing aids help significantly, hearing better in noise is the hearing aid benefit that users agree could be improved further. To hear a conversation better against a noisy background, people intuitively exhibit behaviours that help, regardless of their hearing status. A non-exhaustive list includes head movements, decreasing distance to the desired sound, use of visual and contextual cues, and exploiting familiarity with the conversational partners and knowledge of the language being spoken. The brain has an amazing ability to integrate many channels of information and knowledge to help the person sort out and focus on what they want to listen to. While this ability in and of itself is astonishing, even more astonishing is the efficiency with which it is done. While there is not complete consensus on what the computing speed of the human brain is, a powerful supercomputer would require at least a million times more power to accomplish a similar number of operations per second. This simple fact makes clear how ill-considered it would be to ignore the powerful processing capability of the human brain when designing a hearing enhancement system.

As part of improving the ability of hearing aid users to listen effectively in noise, hearing aids use various noise reduction strategies. The most effective of these so far has been directional microphones which reduce background noise based on the relative spatial locations of speech and competing noise. The signal-to-noise ratio (SNR) benefit for directional microphones in hearing aids can range from 1-2 dB to nearly 10 dB depending on the specific technology and its intended purpose (see Jespersen, 2023[1] for an overview) and the methodology for testing the benefit. For SNR improvement in noise, we are approaching the limits of how effective directionality can be in small ear-level devices, and still be usable in everyday listening scenarios. It has been advised that attenuation of off-axis sound relative to speech from in-front cannot be too large or the user will need to engage in searching to locate and follow sounds of interest.[2] This is effortful and can ironically make listening effectively more difficult. This would work against the goal of using directionality in the first place.

Another aspect related to directional microphones is automatic switching of microphone mode. Switching algorithms make decisions regarding the optimum microphone response depending on the listening environment. These algorithms consider presence of speech and noise in the environment and may include factors such as location of speech or even detected movements of the hearing aid user. Although there may be multiple inputs to the decision-making process, there are few potential modes of operation. Therefore, the automatic behaviour is still relatively simple and easily controlled with a rule-based approach. The Beltone strategy for automatically applying directionality, called the CrossLink Directionality 4 program in premium level products, has 4 microphone modes that are designed to support the natural ways people use their hearing.[3] The rationale prioritises sound quality and spatial hearing, improved SNR and access to surrounding sounds.

Apart from directionality, single-microphone noise reduction – often called digital noise reduction – is used in nearly all hearing aids today. By “single-microphone” it is meant that the sound of interest and the competing noise are collected by the same microphone and therefore mixed in the hearing aid amplification system. A single-microphone noise reduction feature has the potential to reduce background noise even when the sound of interest and the competing noise are co-located. The intention of single-microphone noise reduction is to reduce background noise by determining the spectrum of noise in the input to the hearing aid and reduce gain for only that, leaving the speech intact. Theoretically, this should maintain audibility for speech while reducing noise, thereby improving the SNR. This is easier said than done. To truly work as intended, one needs to know exactly what the spectra of speech and competing noise are at every moment, which is not possible in real life situations. Therefore, a best estimate of these spectra is made based on assumptions that the acoustic properties of speech vary a lot over short time intervals, and that noise backgrounds are typically stable over these intervals. If these assumptions are true, then it is possible to determine when speech is or isn’t present, and to measure the background noise spectrum during pauses of speech. This is exactly the working principle of our legacy single-microphone noise reduction system, Sound Cleaner. Single-microphone noise reduction of this general type can be considered model-based in that it uses prior knowledge of the temporal and spectral features of speech to identify it within the overall signal, while specific knowledge of competing noise is not available.

The consensus so far about the evidence supporting single-microphone noise reduction is that it improves listening comfort,[4] but there is only limited evidence that it affects speech intelligibility.[5] There is also some evidence that noise reduction may reduce working memory load, but not all studies report this finding.[6] Many hearing care professionals may consider single-microphone noise reduction systems to be equivalent across hearing aid brands. However, varying acoustic effects of different systems have been demonstrated, and subjective measures have also revealed, that they lead to different perceptual experiences in both people with normal hearing and hearing loss.[5,7-9]

In fact, large differences in amount of gain reduction and in what frequency regions, activation times, and interaction with different background noises have been demonstrated.[10] The differences in performance and perceptual effects of single-microphone noise reduction systems arise from the many parameters that are available for a single-microphone noise reduction system to adjust. They involve complex calculations of gain reduction in all frequency bands of the hearing aid over short time intervals. How do we know that the right choices are being made to enhance speech and reduce noise at any given moment? The assumptions made in the operation of a model-based system – such as the stability of background noise – may be violated in real listening conditions and could introduce distortions of the speech quality. Therefore, it is relevant to consider whether there could be other ways to make these calculations that could result in even better perceptual outcomes for hearing aid users. As mentioned, the human brain can intuitively apply the right strategies to best achieve this goal. With our Intelligence Augmented approach, our hearing aid noise reduction system can learn to separate and sort incoming sound based on the way humans do.

A data-driven learning approach to noise reduction

Data-driven learning means that the system can analyze large amounts of data to discover patterns and insights without explicitly being taught rules or concepts. To this end, we feed thousands of datapoints in multiple layers that build in complexity for the system to learn from. The advantage this provides over technology based on other methods is that once the algorithm has gathered enough information, it can make its own decisions, without the need for explicit instructions or manual input from a human writing or operating the algorithm. It can better handle fast-changing, novel or ambiguous input because it has previous experience with similar data and can now make informed decisions that are flexible and adaptable.

Our new data-driven noise reduction is called True Noise Reduction. It is built using a common type of AI called a Deep Neural Network (DNN). It gets its name from the neurons in our brain and nervous system. Neurons form the basis for all movements, behaviours, thoughts, and learning in human beings and animals. They enable such amazingly diverse and complex functions because they exist in a large, dense network that shares information quickly and adapts to ever-changing inputs to the system. While each single neuron creates its own reaction from the input it receives, the true power lies in their vast numbers firing together at any given time. A DNN is built to mimic this interconnected neuronal structure. This is done to solve a pre-defined problem by maximising the exchange of information that allows the system to start with small pieces of data and eventually develop its own solution.

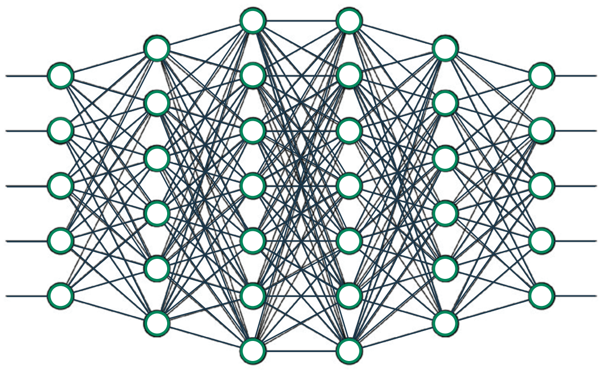

Figure 1. Conceptual diagram of a DNN. Each circle represents one node in the network with each node being analogous to a neuron in the brain. Information flows very quickly across neighbouring nodes.

Balancing performance and efficiency – while keeping the human in control

DNN-based algorithms require vastly more processing power than traditional model-based noise reduction, so much so that a dedicated computer chip is required solely to run the DNN. Hearing aids are already limited in power resources due to their small size, so it becomes even more important to maximise efficiency when using DNN. The architecture and training process were designed with this balance in mind. The greater the number of nodes, and layers of nodes in the DNN, the greater processing resources are required to run the algorithm. The same is true for the sample set of speech and noise used to the train the feature – the more samples used to train; the more power is required from the hearing aid to run it. Therefore, the system was optimised to use as few nodes as possible to achieve contrast between speech and noise and trained using meaningful sound samples commonly found in the real world, including those likely to be considered obvious sources of noise by a human user. To maximise efficiency of speech identification, samples of spoken sentences in various languages and with varied vocal effort were included as target speech. In all, 13.5 million spoken sentences were used in training the Beltone DNN.

There is another vital aspect to this methodology – the human aspect that assumes it is the user's brain, and not the DNN, that makes the most important listening decisions. If the DNN was trained with ambiguous sounds that may or may not be considered noise, not only would power requirements increase, but the feature may begin to "over-process" and reduce gains for sounds or situations where the user may still prefer better audibility. It is the quality of the sound samples, and not the quantity, that make for an effective DNN-based noise reduction.

Efficiency was an explicit design goal of the DNN in True Noise Reduction. Without this, the size of the hearing aid might have needed to be increased to accommodate a larger dedicated DNN chip and battery. Per unit of benefit (quantified as SNR improvement for speech in front and in back in the presence of noise due to the DNN-based noise reduction), True Noise Reduction was calculated to be 17 times more efficient than another leading hearing solution with 4-mic beamforming and DNN noise reduction.[11]

User-selectable True Focus

True Noise Reduction is incorporated in a feature called True Focus for the most effective listening in noise. True Focus combines the efficiency of True Noise Reduction with an upgraded 4-microphone binaural beamformer for better hearing in noise. This feature is the default setting for the Hearing in Noise program and like its predecessor, Ultra Focus 2, is accessible to the hearing aid user via the Beltone HearMax™ app or the onboard hearing aid push button. Figure 2 illustrates the dual-chip design with the settings for True Noise Reduction calculated in parallel to other sound processing with no added delay. Core sound processing from the 360 chip and DNN chip is divided into 17 warped frequency bands, allowing the settings to be easily combined in defining the hearing aid output.

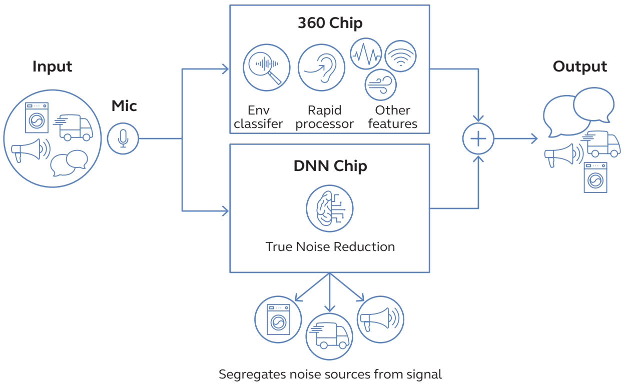

Figure 2. With True Focus, sound enters the hearing aid where SNR improvement for sound in front is achieved via 4-microphone binaural beamforming. With the dual chip design, calculations for True Noise Reduction are carried out by the DNN chip while other sound processing simultaneously is performed by the 360 chip. Settings are combined (at the "+") to define the 17-band warped filter, and then output from the hearing aid.

The upgraded 4-microphone binaural beamformer in True Focus extends directional processing from the 550 Hz used in Ultra Focus 2, down to 250 Hz to optimise the potential for directional benefit. An additional 1 to 2 dB SNR improvement could be expected when the fitting is occluding, and 83% of hearing aid users who participated in a speech recognition in noise task demonstrated better performance with the extended low frequency directional response.[11] Importantly, the unique 3-band structure of the Beltone 4-microphone binaural beamformer continues to be used unchanged in the automatic CrossLink Directionality 4 program to ensure seamless sound quality as microphone modes change.

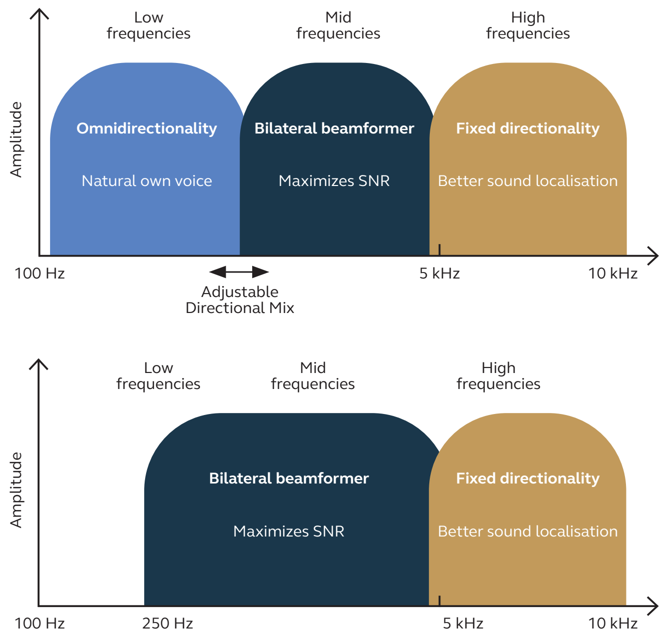

Figure 3. The unique Beltone 3-band binaural beamformer (top panel) provides a large directional benefit while preserving spatial hearing cues and seamless sound quality as microphone modes automatically change in the CrossLink Directionality 4 program. The upgraded binaural beamformer in the user selectable True Focus (bottom panel) extends the low frequency directional response, giving maximum opportunity for SNR improvement.

No compromises

In real life listening, speech often is both the signal of interest and the competing noise. Speech that is the desired sound at one moment can even become noise the next as you change your focus. Even sounds that you might think of as being noise but that inform us about what is happening around us, such as the sound of the refrigerator running or the rustle of someone entering the house, can be important for a hearing aid wearer to hear. An important implication of this is that it is not desirable for any type of noise reduction to completely get rid of what is identified as noise even if it had the power to do so. Eradicating the noise would severely limit access to surrounding sound that would almost certainly interfere with the hearing aid user's listening goals.

Therefore, a primary goal of True Focus with True Noise Reduction has been to provide a natural sound scene with access to both improved SNR for speech in front as well as relief from environmental noise, not elimination of it. The desired effect would spotlight the speech and cause the surrounding noise to recede but still be audible. During development of True Focus, listeners with normal hearing consistently showed strong preferences for the feature over legacy Ultra Focus 2 with model-based Sound Cleaner noise reduction. When tested with hearing aid users, True Focus was preferred 89% of the time over Ultra Focus 2.[11] By designing for efficiency, it has been possible to add the DNN chip to our smallest rechargeable Receiver-In-Ear platform and optimize it with the core hearing aid sound processing of the 360 chip without significantly impacting battery life.

To conclude, True Focus with True Noise Reduction represents a significant advancement in hearing aid technology, leveraging deep learning to enhance sound processing. By combining human cognitive abilities with AI, this innovative approach offers users a more natural auditory experience, paving the way for future developments in hearing solutions.

References

1. Jespersen CT. Hearing aid directional microphone systems for hearing in noise. In Hearing Loss – Diagnosis, Management and Future Challenges, Ciorba, A and Hatzopoulos, S, Editors. London: InTech Open; 2023[cited March 5, 2024]. Available from: https://www.intechopen.com/chapters/1143207

2. Archer-Boyd AW, Holman JA, Brimijoin WO. The minimum monitoring signal-to-noise ratio for off-axis signals and its implications for directional hearing aids. Hearing Research. 2018 Jan 1;357:64-72.

3. Groth J, Kirkwood B, Cui T, Hougaard O, Dieu L. Empowering hearing aid users with directional technology designed for the real world. Beltone white paper. 2023.

4. Bentler R, Wu YH, Kettel J, Hurtig R. Digital noise reduction: Outcomes from laboratory and field studies. International Journal of Audiology. 2008 Jan 1;47(8):447-460.

5. Brons I, Houben R, Dreschler WA. Perceptual effects of noise reduction with respect to personal preference, speech intelligibility, and listening effort. Ear and Hearing. 2013 Jan 1;34(1):29-41.

6. Souza P, Arehart K, Neher T. Working memory and hearing aid processing: Literature findings, future directions, and clinical applications. Frontiers in Psychology. 2015 Dec 16;6:1894.

7. Brons I, Houben R, Dreschler WA. Effects of noise reduction on speech intelligibility, perceived listening effort, and personal preference in hearing-impaired listeners. Trends in Hearing. 2014 Oct 13;18-10.

8. Bentler R, Chiou LK. Digital noise reduction: An overview. Trends in Amplification. 2006 Jun;10(2):67-82.

9. Hoetink AE, Körössy L, Dreschler WA. Classification of steady state gain reduction produced by amplitude modulation based noise reduction in digital hearing aids. International Journal of Audiology. 2009 Jan 1;48(7):444-455.

10. Connor H, Fennen I. Verarbeitungsschwindigheit digitaler Störgeräuschreduzierung. Hörakustik. 2017 May:10-15.

11. Jespersen CT, Dieu L, Rubachandran T. Organic Hearing drives user preference for True Focus. Beltone white paper. 2025.

Worldwide Headquarters
Beltone A/S
Lautrupbjerg 7
DK-2750 Ballerup
Denmark
Tel.: +45 4575 1111
gnhearing.com

CVR no. 55082715

M201690GB-25.01-Rev.A