Hearing Aid Technology to Improve Speech Intelligibility in Noise

Joshua M. Alexander, Ph.D.¹

ABSTRACT

Understanding speech in noise is difficult for individuals with normal hearing and is even more so for individuals with hearing loss. Difficulty understanding speech in noise is one of the primary reasons people seek hearing assistance. Despite amplification, many hearing aid users still struggle to understand speech in noise. In response to this persistent problem, hearing aid manufacturers have invested significantly in developing new solutions. Any solution is not without its tradeoffs, and decisions must be made when optimizing and implementing them. Much of this happens behind the scenes, and casual observers fail to appreciate the nuances of developing new hearing aid technologies. The difficulty of communicating this information to clinicians may hinder the use or the fine-tuning of the various technologies available today. The purpose of this issue of *Seminars in Hearing* is to educate professionals and students in audiology, hearing science, and engineering about different approaches to combat problems related to environmental and wind noise using technologies that include classification, directional microphones, binaural signal processing, beamformers, motion sensors, and machine learning. To accomplish this purpose, some of the top researchers and engineers from the world's largest hearing aid manufacturers agreed to share their unique insights.

KEYWORDS: speech-in-noise, wind noise, classification, directional microphones, binaural signal processing, beamformers, motion sensors, machine learning, hearable, healthable

¹Department of Speech, Language, and Hearing Sciences, Purdue University, West Lafayette, Indiana.

Address for correspondence: Joshua M. Alexander, Ph.D., Department of Speech, Language, and Hearing Sciences, Purdue University, Lyles-Porter Hall, 715 Clinic Drive, West Lafayette, IN 47907-2122 (e-mail: alexan14@Purdue.edu).

Hearing Aid Technology to Improve Speech Intelligibility in Noise; Guest Editor, Joshua M. Alexander, Ph.D. Semin Hear 2021;42:175–185. © 2021. The Author(s).

This is an open access article published by Thieme under the terms of the Creative Commons Attribution-NonDerivative-NonCommercial-License, permitting copying and reproduction so long as the original work is given appropriate credit. Contents may not be used for commercial purposes, or adapted, remixed, transformed or built upon. (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Thieme Medical Publishers, Inc., 333 Seventh Avenue, 18th Floor, New York, NY 10001, USA
DOI: <https://doi.org/10.1055/s-0041-1735174>.
ISSN 0734-0451.

Hearing impairment is often defined by a loss of audibility (the audiogram), and it follows that amplification is the primary means of rehabilitation of that impairment. However, making sounds loud is the easy part—even inexpensive devices can make sound audible for the majority of individuals with mild to moderate hearing loss. As is well-established, the real challenges associated with sensorineural hearing loss include (1) a reduction in the residual dynamic range, which some call loudness recruitment, and (2) distortion at the auditory periphery (e.g., broadened auditory filters), which reduces the signal-to-noise ratio (SNR) of the sensory code before it ascends to the auditory cortex. Unfortunately, the solution to the first of these challenges, wide dynamic range compression (WDRC), can compound the second by reducing spectrotemporal intensity contrasts important for speech understanding by reducing the SNR before the signal even reaches the auditory periphery.¹ Consequently, a lot of care and attention has to go into how and when sounds are amplified. In this respect, one of the most important jobs of a hearing aid is to *control* the amplification process by *attenuating* the right sounds at the right frequencies, times, and direction of arrival. If done correctly, sounds that are meaningful to the hearing aid user, often speech arriving from the front, are preferentially amplified, thereby increasing the SNR and/or reducing listening effort.

Herein lie the challenges faced by hearing aid developers. First, they have to derive solutions to find all of the “right” components in the aforementioned list. However, one developer’s idea of what is “right” might differ from another’s, depending on their philosophies about what is meaningful to the user and assumptions about what the *average* person’s ear and brain need to process the information. Further complicating the matter is that what is meaningful to the user is not always speech and does not always arrive from the front. Therefore, manufacturers’ solutions to the problems encountered by hearing aid users—in the form of algorithms and features—are complex and varied. To most clinicians and academic researchers, the features advertised by manufacturers are often mysterious and viewed anywhere on a scale from mystical to skeptical. However, to make informed decisions

about a manufacturer’s features (e.g., *who* will benefit, *when* they should be activated, *how* to fine-tune them), clinicians must understand how they operate and their limitations. To help fill this need, the content in this issue of *Seminars in Hearing* focuses on hearing aid technology to improve speech intelligibility in noise.

WHO WILL BENEFIT FROM THIS ISSUE?

The intended audience of this issue is broad. First, this issue can be a resource for practicing clinicians to help them get caught up on the latest hearing aid technology. It may seem that a variety of forces are working to make clinicians irrelevant to the overall hearing rehabilitation process. However, the new technology demands that they stay relevant: clinicians are needed to select feature parameters that individualize the hearing aid to a user’s hearing loss, lifestyle, and preferences. Educating current and future clinicians will empower them to make more informed decisions for individual users so that users derive the most benefit from their hearing aids. For future clinicians, the articles in this issue will help fill the void of resources on current hearing aid technologies available to audiology graduate students. This open-access issue can supplement current textbooks on hearing aids, which are often targeted toward beginners and more tried-and-true technologies.

In the spirit of educating the reader, authors were asked to provide the following: (1) an introduction to the problem their selected technology addresses; (2) the nuances involved in finding technological solutions to the problem; (3) general signal processing solutions that have been proposed in the past; (4) how their signal processing approach works and what makes it innovative; and (5) data to support the efficacy of their technology. Given this outline, these articles also can be a resource for curious engineering students and hearing scientists interested in understanding the technological and clinical challenges that research needs to address. For this audience, information is provided to help them understand the human or perceptual constraints that inform decisions that need to be made when developing a signal processing solution. Finally, while the fine details about the

solutions (e.g., equations) are not included in the issue, the authors have provided a rich set of references to support those interested in the engineering behind them.

The ambitious goals set forth in the preceding paragraphs could not be achieved without the invaluable contribution of time and talent by the individual authors, who represent the top researchers and engineers in the hearing aid industry, and their employers, who allowed them to take time away from other projects and who provided the financial support to make each article in this issue open access so that it can reach as many people as possible. The willingness of those in the hearing aid industry to contribute in so many ways to this issue shows they value the role of clinicians in the provision of hearing aids and understand the need to educate clinicians about why their technology solves a problem that clinicians may not have known existed.

EDITORIAL DECISIONS: MORE THAN MEETS THE EYE

I cannot thank the authors enough for their contributions. I especially thank the authors for allowing me to deliberately orchestrate this issue: first, by selecting the topics I wanted them to write about; then, by liberally editing their drafts to make them “sing” together, that is, to bring uniformity and clarity across the integrated whole. Much to the authors’ chagrin, I changed their terminology to promote consistency throughout the articles. This was easier said than done. For example, a decision had to be made about what to call the things that individuals with hearing impairment put in their ears to help them hear: aids, instruments, or devices. I opted for “hearing aids” for historical continuity, especially with the research literature. I say this while acknowledging that the term “hearing devices” is becoming more popular since it more aptly encompasses the broad range of functions manufacturers are putting in today’s hearing aids. At a basic level, a person with hearing loss can use hearing aids in the same way that a person with normal hearing uses their wireless earbuds to stream music, phone calls, television, navigation systems, etc., and to send voice commands to their smartphones as when

interacting with a virtual assistant. These features fall under the category some call “hearable” technology. However, as discussed by Fabry and Bhowmik, some hearing aids are now using embedded sensors that allow them to function as a “healthable” technology. For now, I will stick with the term “hearing aids” since this is their primary function and continues to be the reason why people buy them.

Along similar lines, a decision had to be made about what to call the people who put these things into their ears to help them hear: patients, wearers, users, or listeners. The term “patients” is too inclusive for the broad range of service delivery models available today. The term “wearers” works if the goal is for individuals to put the hearing aids on and more or less forget about them, much like individuals who wear glasses. However, many of the technologies designed to help people hear better in noise allow them to override the automatic feature selections so that they can be activated, deactivated, or adjusted in some way to accommodate a preference or different intent for listening. For these reasons, I opted for the term “user.” The term “listeners” is sometimes used in these articles to refer to the participants in a laboratory experiment, especially since not all of the experiments involved hearing aids and not all of the participants used hearing aids outside of the laboratory.

The last term related to the provision of hearing aids is the person providing the hearing aids and completing the necessary rehabilitative care (fitting, counseling, etc.): audiologist, clinician, dispenser, or health care professional (HCP). As a professor who teaches students seeking a Doctorate of Audiology (AuD), “audiologist” and “clinician” are my preferred terms. The term “dispenser,” in my opinion, is too sterile and fails to fully capture the caring aspect of the relationship and the professional training involved. In some places, like France and Québec, audiologists are not allowed to fit hearing aids; instead, this is performed by highly trained “hearing aid acousticians” (audioprothésistes). Perhaps, for these reasons, and in the spirit of inclusivity, some authors opted for the term “hearing care professional.” Given my academic bias, while striving to remain inclusive, I opted for a happy medium with the use of the term “clinician.” The reader

is free to substitute whatever term best fits their service delivery model.

Finally, a frequent term that the reader will come across is “listening environment,” which includes all of the acoustic and nonacoustic factors that can influence a person’s communicative or noncommunicative intent when using their hearing aids in a particular time and place. I had the hardest time convincing authors to use this term. Other terms favored by different authors included “acoustic ecology,” “sound environment,” “auditory reality,” and “situation.” While all or most information used by a hearing aid is acoustic, as discussed by Branda and Wurzbacher, some hearing aids now include information about the user’s motion. In addition, I like the term “listening” because it implies that the individual user’s goals in a particular environment need to be considered when deciding which sounds, frequencies, times, and direction of arrival to amplify and which to attenuate. I credit Donald Hayes for distinguishing between a hearing aid user’s real-world experiences in different listening environments and contrived research setups that try to mimic these listening environments, for which he used the term “acoustic scene.”

ORGANIZATION OF TOPICS

While all of the aforementioned terms may seem like minor details, my end goal was to make it easier for the reader to move from one article to

another and draw connections between them without slowing down to decide if the terms refer to the same concept. In addition to the use of common terminology, topics were carefully selected and organized to cover the breadth of technologies available in today’s hearing aids while minimizing overlap in content between articles. This being said, the different technologies (solutions) contained in a single hearing aid are necessarily integrated, which makes it challenging to write about them in isolation. Therefore, each article contains several citations to other articles in this issue so that wherever readers start, they might end up going through the entire issue cover-to-cover. As shown in Fig. 1, the interconnectedness between the articles centers on four main themes: automatics, directionality, noise reduction, and artificial intelligence (AI).

Automatics

To varying extents, concepts related to automatic program selection and feature adjustments are woven into every article. This is probably a reflection of the fact that one of the goals behind modern hearing aid design is the ability to react, with or without user intervention, to changing listening environments to optimize the balance between speech intelligibility and listening comfort or environmental awareness. WDRC is one of the earliest and most primitive forms of automatic processing. WDRC works well when the signal is speech in

Figure 1 Venn diagram showing the broad topics that are the primary focus of the articles in this issue. Numbers correspond to the different articles listed in the legend.

quiet, but it breaks down as the complexity of the environment increases. To make informed decisions about how to alter gain, directionality, noise reduction, etc., a hearing aid needs an accurate account of the listening environment: overall level, SNR, type of noise, relative positions of different sound sources, etc. It is the job of the environmental classifier to compile this information and decide what type of listening environment a user is in. Program and feature adjustments are then performed based on the philosophies of the hearing aid developer and their assumptions about the user's listening intent (e.g., hear a conversation, listen to music, block out background noise).

Hence, in many respects, classification systems are the bedrock of modern hearing aids. For this reason, the article by Donald Hayes on environmental classification leads this issue of *Seminars in Hearing*. As highlighted by Hayes, the accuracy of the environmental classifier is so critical because if it makes a mistake, the subsequent adjustments by the hearing aid can ruin a perfectly good fitting and cause the user to struggle more when communicating. Given the importance of the environmental classifier, one would think that different manufacturers would be in good agreement. Hayes reports on an elaborate study conducted in collaboration with David Eddins and Erol Ozmeral at the University of South Florida. Together, they created a multitude of acoustic scenes to simulate a range of listening environments varying in complexity. Then, they compared the datalogging results from five brands of hearing aids from different manufacturers to the classifications from a group of normal-hearing listeners. It may or may not surprise the reader to learn that the five hearing aid brands had a high agreement for the easiest acoustic scene—speech in quiet—but diverged significantly as the acoustic scenes increased in complexity. These results might be explained by the number of acoustic factors a hearing aid developer has to consider when designing a classification system.

Finally, Hayes reports on the Global Listening Environment Study. He collected information regarding the proportion of time a large, worldwide group of hearing aid users spent in different listening environments as recorded from their hearing aid classifier's raw output

(moment-by-moment probabilities). Their primary findings were (1) users spent the most time in quiet, followed by a conversation in a small group, and then by a conversation in quiet; the smallest amount of time involved noise, crowds, and music; (2) this pattern was consistent regardless of how the participant demographics were broken down (e.g., gender, age, population density); and (3) within each demographic, there was substantial individual variability, such that individuals had large deviations from this pattern.

Interestingly, the other three articles in the automatics category in Fig. 1 (Branda and Wurzbacher; Balling et al; Fabry and Bhowmik) discuss technologies designed to deal with situations where the acoustic information traditionally compiled by the hearing aid fails to accurately capture the listening environment and/or the user's listening intent. The article by Eric Branda and Tobias Wurzbacher discusses how motion sensors (accelerometers) can help the hearing aid classify the listening environment more accurately. In particular, motion sensors provide information about the user's movement within the listening environment, indicating that their listening needs may have changed. Just as the amount of gain provided by a hearing aid must balance audibility and listening comfort, the amount of directionality must balance focus and environmental awareness. As discussed in the section on directionality, different manufacturers have different philosophies and approaches to address this challenge. The most basic solution is to lessen the amount of directionality or switch to omnidirectional mode when the hearing aid identifies the listening environment as one where the user will likely benefit from environmental awareness. More particularly, Branda and Wurzbacher identify the user's movement within the listening environment (walking) as a key indicator that they will benefit from having auditory access to information arriving from all directions. By way of example, they describe a busy restaurant as a listening environment where two individuals can be exposed to the same acoustics but have different listening intents: (1) a patron sitting at a table facing their conversation partner would benefit from some amount of directionality that favors sounds arriving from the front; (2) a waiter

walking in the same listening environment would benefit from full environmental awareness.

The remaining two articles in the Automatics category in Fig. 1 (Balling et al; Fabry and Bhowmik) discuss technologies related to a class of automatics based on AI. They discuss user-initiated technologies designed to circumvent problems that arise when automatic program settings and feature adjustments fail to meet an individual user's needs at a particular moment and time. Laura Winther Balling and colleagues identify several reasons why the default or automatic settings in a hearing aid may be suboptimal for an individual user, place, and/or time. First, general problems can arise from the simple fact that hearing prescriptions and automatic solutions programmed into the hearing aid are designed for an average user. The programming decisions may be based on actual data or assumptions made by the developers of the algorithms that drive the automatic processing. Second, specific or local problems can arise when the automatic adjustments do not match the user's listening intent in a particular listening environment.

Like, Branda and Wurzbacher, Balling et al describe how two individuals in the same listening environment can have different listening intents, each requiring different signal processing solutions. Balling et al provide examples of static listening environments composed of multiple sound sources, any of which could be the target of the user's attention. As highlighted by Balling et al, customized solutions for a particular listening environment can be created by simple gain adjustments in three frequency channels; however, even this would require an enormously large number of comparisons. Therefore, the core of their technology uses machine learning to iteratively refine a series of A-B comparisons so that an optimal solution can often be found using 20 or fewer comparisons. When users desire a change in their hearing aid settings, they activate an app on their smartphone which interfaces with their hearing aids. Part of what makes the technology described by Balling et al so efficient is that it gathers the degree of preference for one setting over another, which provides the machine-learning algorithm with a continuous range of

values rather than a discrete set of responses. The technology is also valuable from a research and development perspective because it yields data about users' preferences in many different listening environments.

As indicated by David Fabry and Achintya Bhowmik, most classification systems, including those discussed in the article by Hayes, are based on solutions derived from AI. Machine learning approaches have been used to identify patterns of acoustic features that reliably classify sounds into a small number of categories related to different listening environments. Despite extensive training on large databases of acoustic scenes, Fabry and Bhowmik indicate that the accuracy of most classification systems tops out at around 80 to 90%. For the remaining times when the pre-programmed, automatic solution fails to meet users' listening needs, they can initiate a change in their hearing aid settings by double-tapping MEMS-based (micro-electro-mechanical systems) motion sensors on their hearing aids. This user-initiated action causes the hearing aid to capture an acoustic snapshot of the listening environment which is then analyzed using a form of processing known as "edge computing." Edge computing is an emerging area in industrial applications that most often involve one or more sensors (e.g., wind speed, fluid pressure) whose output is efficiently analyzed in real-time by a local device connected to the sensor. Because no data transfer to another computing device or network is required, automatic decisions or actions can be made with minimal delay. The acoustic snapshot is processed within the hearing aid by an onboard AI model trained with machine-learning technology. Subsequently, the parameters for eight different preset classifications are dynamically manipulated to optimize speech intelligibility and sound quality. Some of the parameters affected include settings for gain, output, noise management, and directional microphones.

To demonstrate the versatility of onboard, on-demand edge computing within the hearing aid, Fabry and Bhowmik describe the communication challenges brought about by the use of face masks during the COVID-19 era. It is well-documented that people, including those with normal hearing, benefit tremendously from being able to see a talker's mouth. Because

opaque face masks remove this benefit, individuals need to rely more on the acoustic signal. Face masks can compound the problem by reducing high-frequency speech energy. Unfortunately, face masks and shields with a clear window that allows one to see a talker's mouth tend to alter the speech acoustics the most. So, individuals with hearing loss are at a disadvantage either way. A simple solution would be to compensate for the altered speech acoustics introduced by a face mask via a simple gain adjustment in the hearing aid; however, different styles of face masks have different effects on the acoustics. A more sophisticated solution described by Fabry and Bhowmik benefits from edge computing because adjustments to several features can be made regardless of the face mask style, the distance between conversation partners, and the presence of background noise.

Directionality

BINAURAL PROCESSING

In addition to frequency, intensity, and time (phase), direction of arrival is the fourth dimension of sound that provides organisms, including people, with a rich source of information about the auditory scene within their listening environment. It follows that direction of arrival is also a rich source of information hearing aids can use for classification and signal modification. In both people and hearing aids, the sense of direction is greatly enhanced when comparisons are made between two ears or sets of microphones positioned on opposite sides of the head. Peter Derleth and colleagues provide a comprehensive overview of the cues used for binaural hearing and the perceptual benefits of binaural hearing, especially when listening to speech in noise. Derleth et al indicate that hearing aids can distort a user's spatial processing via (1) acoustic coupling, (2) independent signal processing, and (3) beamforming.

Acoustic coupling refers to the relative position of the hearing aid microphones and the degree to which the fit is opened/occluded. Monaural cues critical for localization in the vertical plane (up-down; front-back) are completely lost when microphones are placed above the pinna because the user no longer has

access to their unique pinna acoustic filtering properties. Furthermore, as the fit becomes more occluded, users lose direct access to low-frequency binaural cues (interaural time differences, ITDs) critical for localization in the horizontal plane (left-right), forcing them to depend on indirect access to these cues via amplification.

Binaural cues also can be distorted by independent signal processing in the left and right hearing aids. While ITDs can be affected, interaural level differences (ILDs) are most susceptible since the majority of hearing aid algorithms manipulate level, including WDRC, noise reduction, and directionality. According to Derleth et al, synchronizing the signal processing behavior in each hearing aid can help preserve binaural cues, but this might limit the overall effectiveness of the algorithm. For example, the SNR improvement from a noise-canceling algorithm will likely tradeoff with binaural cue preservation.

Despite its name, binaural beamforming can significantly distort a user's spatial processing because it wirelessly combines the signals from the microphones on the left and right hearing aids into one highly directional signal in the output. Because the left and the right ear receive the same diotic signal, binaural cues which rely on a comparison between ears (ITDs and ILDs) are completely lost. For this and the previously described reasons, Derleth et al note that the challenge faced by developers of binaural hearing aid algorithms is to provide the acoustic cues necessary for a user to create and maintain an individualized, mental spatial map while providing them with the ability to focus their attention on a single acoustic source if needed. In so doing, the likelihood that an individual hearing aid user will benefit must be considered, which will depend on their residual auditory capabilities, listening needs, and listening intent. One way of meeting this challenge is to limit beamforming above a certain frequency range. For individuals with relatively normal low-frequency hearing, this can easily be achieved with an open acoustic coupling, which will naturally preserve the ITDs in the low frequencies. As noted by Derleth et al, this approach works because the low-frequency ITDs dominate binaural

perception if the high-frequency ILDs provide contradictory information. For individuals who need amplification in the low frequencies, an omnidirectional response pattern can be applied in that frequency range. Results presented in their article indicate that this approach is a good compromise between preserving and ignoring binaural cues and provides the best speech intelligibility regardless of the noise scenario and the listeners' performance on a perceptual test of binaural sensitivity.

DIRECTIONAL MICROPHONES

As noted in the "Introduction," one of the primary challenges associated with sensorineural hearing loss is the reduced SNR in the sensory code. It follows that amplification by itself will not alleviate this problem. One tried-and-true technology for combating the problem is directional microphones which enhance the SNR of the acoustic signal before it reaches the cochlea. Continuing with the theme of controlling amplification, directional microphones increase the SNR by attenuating sounds arriving from directions that are not occupied by the signal of interest. As discussed by Charlotte T. Jespersen and colleagues, despite their potential to improve speech intelligibility in noise, the SNR benefit provided by directional microphones is contingent on certain conditions. The speech source must be (1) spatially separated from the noise sources; (2) positioned in a direction that will not be attenuated by the directional microphones, usually in front of the user; and (3) relatively close to the user. Not surprisingly, these conditions are not always met in a user's listening environment; hence, they do not always benefit from directionality.

Depending on the aforementioned conditions, the listening environment, and the user's listening intent, users often desire to have an awareness of sounds arriving from all directions. This awareness allows users to hear if someone, who is not in their field of view, is trying to get their attention (e.g., the waiter in Branda and Wurzbacher's article). It also can help promote a sense of psychological comfort via the increased environmental vigilance afforded by "surround sound." As noted by Derleth et al, one challenge faced by hearing aid developers is to provide users with the acoustic cues necessary

for binaural processing, including environmental awareness, as well as provide them with the ability to focus their attention on a single acoustic source.

One low-tech solution to the above challenge is a user-initiated switch in the microphone mode of one or both hearing aids. Alternatively, Jespersen et al discuss a high-tech, automatic solution that coordinates the microphone mode of each hearing aid. First, the directional microphones on each hearing aid are used to assist environmental classification by providing information about the relative locations of the speech and noise sources. Then, 2.4-GHz wireless technology is used between the hearing aids to coordinate their microphone modes to create one of three bilateral modes. One bilateral mode is designed to promote environmental awareness in quiet or moderately complex listening environments that may or may not have speech. It does this by (1) using omnidirectional responses in both hearing aids; (2) synchronizing WDRC behavior to preserve the high-frequency ILDs; and (3) simulating or preserving pinna cues. For hearing aids where the microphone is above the pinna, pinna cues are simulated by giving the high frequencies a forward-facing directional response. Natural pinna cues also can be preserved by putting the microphone directly in the ear canal using a new hearing aid style called Microphone and Receiver-In-the-Ear (M&RIE).

Another bilateral mode is designed for listening environments where the user will benefit from both environmental awareness and enhanced SNR. This mode is triggered when the noise arriving behind the user exceeds a certain level threshold. Speech may be present, but not solely in front of the user. One hearing aid has a directional response. The side that is chosen depends on the relative positions of the speech and noise. The other hearing aid has a specifically designed omnidirectional mode that attempts to compensate for the effect of the head shadow on its sensitivity to sounds arriving from the opposite side of the head.

The last bilateral mode is designed for noisy listening environments where speech is detected from the front only. In this case, a weighted binaural beamformer is used. In listening environments where the noise is diffuse,

the beamforming algorithm weights the inputs of both hearing aids' microphones equally. However, in listening environments with more noise on one side or the other, the algorithm takes advantage of the head shadow by assigning greater weights to the side with less noise than the side with more noise. Furthermore, unlike other beamformers that use a two-band system with low-frequency omnidirectional processing and high-frequency beamforming (e.g., Derleth et al), this beamformer uses a three-band system whereby frequencies above 5,000 Hz are processed with a monaural hypercardioid response to help preserve the ILDs in this frequency range.

Noise Reduction

The first of the two articles that specifically discuss noise reduction involves a problem—wind noise—which is exacerbated by the way most directional microphones operate. While seemingly benign compared with all the other challenges faced by hearing aid users, wind noise can significantly reduce the usability of hearing aids in the outdoors due to its effects on speech intelligibility and listening comfort. The article by Petri Korhonen discusses the subtle but complex properties of wind and movement that create problems for hearing aid users and hearing aid developers. Korhonen describes wind noise as an artifact caused by induced random pressure fluctuations near a microphone's diaphragm. It is distinct from the environmental noise that wind might directly or indirectly cause when airflow encounters an obstruction. Hence, conventional single-channel noise reduction techniques are less effective at combating the side effects caused by wind.

Directional microphones increase the adverse effects of wind noise. Directional microphones delay the output from the rear microphone and subtract it from the output of the front microphone. The more correlated or in phase the sound is at the two microphones, the more effective the cancellation. Low frequencies have longer wavelengths which cause them to be more in phase at the two microphones, which causes their level to drop significantly. To compensate for this low-frequency roll-off, directional microphones often have an

equalization filter that boosts the low frequencies. Because wind causes random (uncorrelated) pressure fluctuations at each microphone, the outputs from the two microphones add instead of subtract. Furthermore, the assumed low-frequency roll-off is absent, so the equalization filter makes the low-frequency wind noise even more intense.

Korhonen discusses several low-tech solutions to combat wind noise, including auto-switching to an omnidirectional microphone mode in at least the low frequencies or adaptively reducing the gain in the low-frequency channels. In addition, the hearing aid can add, instead of subtract, the output of the two microphones, which will increase the SNR of the correlated far-field sounds relative to the uncorrelated environmental and wind noise sources. It also is standard practice to use a cover, shield, or foam on or around the microphone diaphragm to laminate the airflow or reduce its velocity. Finally, positioning the hearing aid microphones in the canal, including the M&RIE receiver discussed by Jespersen et al, or in the folds of the pinna provide natural relief from wind noise.

Korhonen also discusses several high-tech solutions to combat wind noise. One technique involves binaural streaming of the low-frequency part of the microphone output from the hearing aid on the side of the head that is less exposed to the effects of wind noise to the hearing aid on the opposite side. According to Korhonen, algorithms trained using machine learning may show promise but have yet to demonstrate significant benefits. The wind noise attenuation algorithm discussed by Korhonen employs an adaptive filter similar to those used for conventional single-channel noise reduction (e.g., a Wiener filter). Instead of viewing the directional microphones as a hindrance, the approach exploits the signals at the two microphones to reduce the wind noise level. In basic terms, the algorithm uses an adaptive filter to model the differences between the outputs from the two microphones, thereby separating the correlated and uncorrelated parts of the signal. The SNR can be increased because the correlated part corresponds to the environmental signal, including speech, and the uncorrelated part corresponds to the wind noise.

The remaining article in this overall summary combines elements of directionality, noise reduction, and AI. However, as indicated by Asger Andersen and colleagues, directionality (what they refer to as “beamforming”) can be considered one component of noise reduction, especially when used in conjunction with “post-filtering” techniques that many think about when they refer to “noise reduction.” Conventional single-channel noise reduction techniques provide users with listening comfort in noise, but they cannot improve speech intelligibility. In fact, they are often prevented from being too aggressive so that speech intelligibility is not adversely affected.

Andersen et al describe a technique for improving speech intelligibility in noise that integrates directionality with postfiltering. First, adaptive directionality is applied to the output of a filterbank to optimally attenuate the different frequency components of noise that arrive from a different direction as the speech. Then, post-filtering with an adaptive filter (e.g., a Wiener filter) is used to attenuate noise sources arriving from directions near the speech source. To know which frequencies to attenuate at which times, Wiener filters use the statistical properties of speech and noise to estimate the short-term SNR in each frequency band. By integrating directionality with postfiltering, the directional system can be exploited to derive a more accurate estimate of SNR. This is possible because one directional pattern can be presented to the user while the hearing aid simultaneously evaluates other directional patterns for use by other algorithms. The increased accuracy of the SNR estimates provided by the integrated system can increase speech intelligibility, unlike conventional single-channel noise reduction techniques.

Andersen et al describe an even more sophisticated technique for improving speech intelligibility in noise whereby directionality is integrated with a noise reduction algorithm that employs deep learning. Traditional post-filtering techniques rely on relatively simple models based on statistical properties the hearing aid developer thinks will help separate speech from noise. However, the real speech-in-noise problem is far more intricate than a human can model mathematically. For this, machine learning can be used to find more sophisticated

solutions. The technique described by Andersen et al uses a deep neural network (DNN) to find a way to make examples of noisy speech that were processed by a hearing aid to be similar to their clean speech counterparts. The resulting algorithm is free to model whatever structures can be discovered in examples, which are likely to be mathematically complex and difficult to explain. The algorithm’s accuracy is only as good as the number, variety, and realism of the examples used to train the DNN. For this, the authors describe an elaborate method of obtaining multiple real-world examples of noisy listening environments using a spherical array of 32 microphones. Thus, the noisy listening environments could be rendered in a sound studio with a high degree of dimensionality using an equal number of loudspeakers. Speech examples were obtained from talkers who listened to the recordings of the noisy listening environments under headphones. This way, they had recordings of the clean speech that was uncontaminated by the noise. Finally, convolving the speech and noise recordings with impulse responses from different people’s ears and different hearing aid styles yielded a vast database of examples used to train the neural network.

Andersen et al evaluated the efficacy of the new algorithm using listeners who were tested on a hearing aid with the DNN-based postfiltering and a hearing aid with traditional postfiltering. Both hearing aids had the option to integrate postfiltering with directionality. First, with directionality deactivated, they documented objective improvements in SNR for the DNN-based postfiltering technique versus the traditional postfiltering technique. SNR improved even more for both techniques when directionality was activated, especially for the hearing aid with traditional postfiltering, thereby modestly decreasing the advantage of DNN-based postfiltering. This result suggests that DNN-based postfiltering without spatial information shares some properties with directionality for segregating speech from noise. Consequently, behavioral results indicated that DNN-based postfiltering alone significantly improved speech intelligibility more than the control condition (postfiltering and directionality deactivated). DNN-based postfiltering alone also resulted in higher speech intelligibility than traditional postfiltering alone,

which was not significantly different from the control condition. Thus, unlike conventional noise reduction techniques, DNN-based postfiltering improved speech intelligibility by separating speech from noise (albeit not as much as directionality). Finally, Andersen et al present data that suggest auditory scene segregation was enhanced by the hearing aid with integrated DNN-based postfiltering and directionality. Compared with the control condition and the hearing aid with traditional postfiltering, pupilometry and electroencephalography indicated improvements in listening effort, selective attention, and divided attention.

Artificial Intelligence

This summary comes full circle with a brief discussion on AI. As discussed by Andersen et al, Balling et al, and Fabry and Bhowmik, AI is a relatively broad term that encompasses many methods and processes, including machine learning and two of its subfields, Bayesian optimization, and deep learning. Increased computing power, knowledge, and awareness have allowed AI to permeate many aspects of our everyday lives, and hearing aids are no exception. For example, AI is used in system design, including the development of classification systems (Hayes; Fabry and Bhowmik), wind noise reduction algorithms (Korhonen), and postfiltering algorithms (Andersen et al). Furthermore, there are now AI-based, user-initiated applications onboard the hearing aid (Fabry and Bhowmik) and a smartphone (Balling et al). AI is even being used on large datasets to find patterns that can help improve existing algorithms and inform decisions about hearing aid design (Balling et al). Finally, in conjunction with embedded sensors, AI is expanding the role of hearing aids in promoting overall health, such as automatic fall detection and health monitoring (Fabry and Bhowmik).

CONCLUSION

The topics covered in this issue were selected to represent technologies used generally and specifically in the hearing aid industry. Authors were invited to write about specific topics;

therefore, their contributions should not be interpreted as representing the manufacturers' latest and greatest technology for improving speech intelligibility in noise. Upon reading this issue, readers are encouraged to connect the material with information gathered from the manufacturers they are most familiar with. Furthermore, the topics in this issue are not inclusive of all the technologies available to help hearing aid users understand speech better in noise. For example, wireless technology, including remote microphones and smartphone applications, is briefly discussed in this issue by Fabry and Bhowmik. A more comprehensive review of these topics is the subject of previous issues of *Seminars in Hearing*: Volume 35, Issue 3 (2014) and Volume 41, Issue 4 (2020). Finally, with all of the technologies currently available for improving speech intelligibility in noise and with those currently in development, it is clear that this is an exciting time to develop, research, fit, and use hearing aids.

CONFLICT OF INTEREST

None declared.

ACKNOWLEDGMENTS

This issue would not have been possible without the invaluable contributions of Donald Hayes, Peter Derleth, Eleftheria Georganti, Matthias Latzel, Gilles Courtois, Markus Hofbauer, Juliane Raether, Volker Kuehnelt, Charlotte T. Jespersen, Brent C. Kirkwood, Jennifer Groth, Eric Branda, Tobias Wurzbacher, Petri Korhonen, Asger Heidemann Andersen, Sébastien Santurette, Michael Syskind Pedersen, Emina Alickovic, Lorenz Fiedler, Jesper Jensen, Thomas Behrens, Laura Winther Balling, Lasse Lohilahti Mølgaard, Oliver Townend, Jens Brehm Bagger Nielsen, David A. Fabry, and Achintya K. Bhowmik.

REFERENCE

1. Alexander JM, Masterson K. Effects of WDRC release time and number of channels on output SNR and speech recognition. *Ear Hear* 2015;36(02):e35–e49