

SUMMER 2025

Exploring the usage of AI in hearing aids to improve their functionality

By Calvin Zou

AUTHOR BIOGRAPHY

Calvin Zou, a 17-year-old student at Watchung Hills Regional High School in New Jersey, aspires to pursue a career in the STEM field, with a particular interest in data analytics. Inspired by the rapid advancements in artificial intelligence, Calvin is especially intrigued by its potential to enhance hearing aid technology. His interest is deeply personal—his great-grandmother struggles with severe hearing loss and finds it difficult to hear him during visits, even with the aid of current hearing devices. Living in a noisy urban environment only adds to the challenge. Calvin hopes that continued progress in AI will one day allow his great-grandmother to clearly hear his voice again.

ABSTRACT

Hearing loss currently affects around 430 million people worldwide and this number is expected to climb even higher in the future, reaching up to 700 million in 2050. With current hearing aid technology falling short in addressing complex auditory challenges, artificial intelligence (AI), particularly through machine learning, offers promising solutions to improve human speech clarity. By utilizing machine learning's pattern recognition capabilities in complex noise environments (environments with multiple sounds), AI can minimize noise, which improves speech clarity. Additionally, AI has the potential to optimize hearing aids to fit a user's sound profile, leading to better speech enhancement. Its potential to reduce the need for frequent audiologist visits could also significantly lower costs. In this study, a framework for how AI can be utilized to mitigate a complex noise environment is explored. These feasible advancements make AI a potential catalyst in revolutionizing the field of hearing.

Keywords: *artificial intelligence, machine learning, hearing aids, audiology, active noise cancelling, complex noise, pattern recognition, deep neural network, Starkey*

SUMMER 2025

Artificial Intelligence (AI) has been a huge topic as of late, and it has major potential to revolutionize the hearing industry. So what is AI, exactly, and what can it do for hearing?

The goal of AI is to simulate human intelligence. AI can do this through visual reasoning, abductive reasoning, or deductive reasoning. This paper explores how machine learning, a subset of AI, can be used to simulate human intelligence for audio processing. Machine learning is a method by which AI can learn to make predictions through large datasets, essentially mirroring how the human brain recognizes patterns.

Hearing loss ranks as the third most prevalent chronic physical condition, affecting 430 million worldwide. This number is expected to climb even higher in the future, up to 700 million in 2050 (World Health Organization, 2025). In the U.S., hearing loss affects 15% of American adults, roughly 37.5 million people (Victory, 2024). There are two types of hearing loss: conductive and sensorineural. Conductive hearing loss happens from infections, earwax, or abnormalities in the ear's structure. This type of hearing loss is due to an obstruction. The other type of hearing loss, sensorineural, happens from nerve damage in the inner ear (John Hopkins Medicine, 2025). This includes genetic conditions, age, or prolonged exposure to loud noises. Sensorineural hearing loss is the type that typically requires hearing aids (John Hopkins Medicine, 2025).

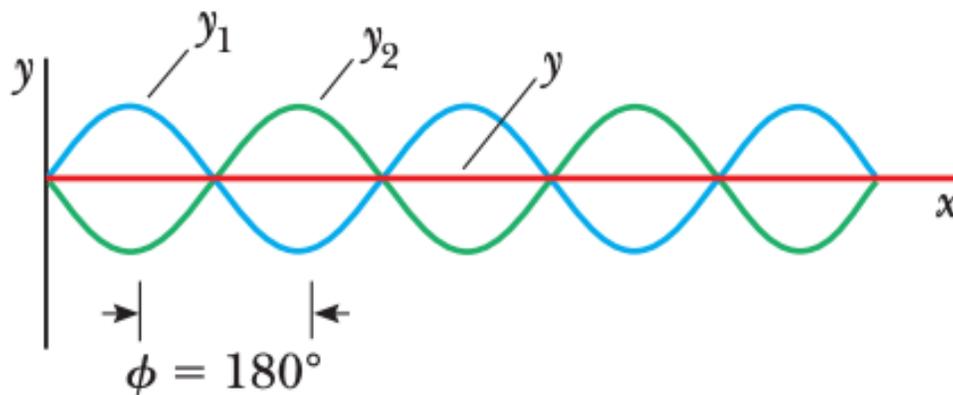
In the U.S, hearing aids are used by 6.4 million Americans, with an estimated 28.8 million more who could benefit from using them (Victory, 2024). One of the two most common complaints about hearing aids is their cost; prices typically range from \$2,000 to \$7,000, with lower-cost models often providing inferior quality compared to their higher-end counterparts (Mroz, 2024). Hearing aids also necessitate multiple clinic visits for testing, fitting, and ongoing monitoring, further contributing to overall costs. In New Jersey, the average cost of these visits ranges from \$96 to \$142 (Sidecar, n.d.). The second major concern is the current limitation of hearing aids in effectively managing noise, which is defined as unwanted or excessive sound (Thomas, 2024). Traditional hearing aids rely on basic amplification, which increases the volume of all sounds, including background noise, making it difficult for users to focus on the desired sounds, such as a conversation. In crowded settings with multiple sounds, hearing aids not only amplify just the speech but also all surrounding sounds. This leads to overwhelming noise levels that make it hard for users to discern specific conversations.

Many applications of AI are still in the prototype stage, but the implementation of AI-what it can do in the hearing space-has been theorized already. A primary challenge is background noise. In a perfect scenario, AI can tailor the sound quality of the hearing aids and use noise cancellation to drown out the background noise. This results in the patient with only the sound of the person they want to hear (Sygrove, 2024). A secondary concern is sound quality, which can be improved through the application of AI. Lastly, the issue of cost can also be diminished, as the need for multiple visits to adjust and test is lessened.

NOISE CANCELLATION

Noise cancellation refers to the use of technology to cancel out external noise. It is traditionally used in earbuds and headphones. While its application in hearing aids has been explored to some extent, the recent surge in AI advancements has driven more in-depth research and practical implementation of noise cancellation in this field. Active Noise Cancellation (ANC) involves the process to reduce/eliminate unwanted background noise. ANC utilizes destructive interference by generating opposing sound waves to cancel out the unwanted sound (Thomas, 2024). ANC operates through a three-step process. First, external microphones capture ambient noise (Thomas, 2024). Second, the system analyzes the incoming sound waves, identifying their amplitude and phase metrics (Thomas, 2024). In the final step, the system generates an anti-noise signal with the same amplitude but an inverted phase relative to the original sound wave (Fig. 1), effectively canceling the unwanted noise through destructive interference (Thomas, 2024).

Figure 1
Sound wave and counter wave



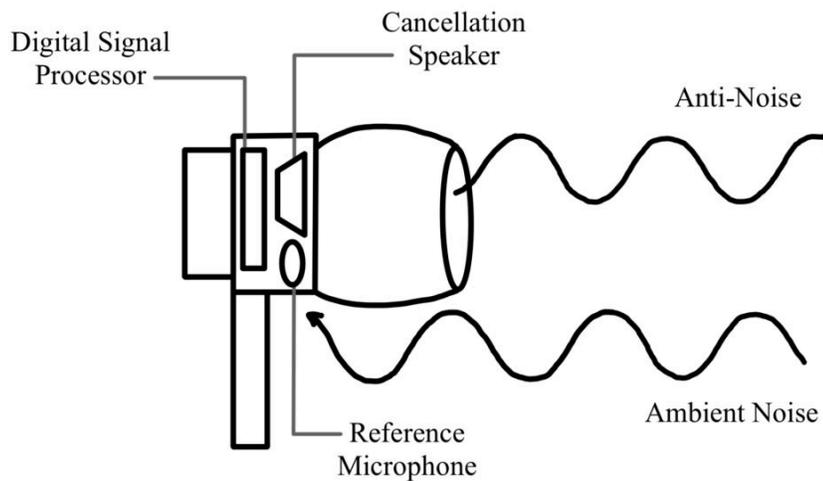
Note. This graph depicts the counter wave and the original sound wave, where y_1 is the original and y_2 is the counter. Notice the amplitude being the same but the phase being inverted. Reprinted from “Engineering Healthy Silence: Using Noise-Cancelling Headphones to Block Harmful Sound” by R. Pak, 2024 (<https://illumin.usc.edu/engineering-healthy-silence-using-noise-cancelling-headphones-to-block-harmful-sound/>).

As noise is never a single sound wave but rather a continuous stream of waves, modern ANC devices employ real-time adjustment to dynamically cancel persisting sounds such as an engine hum.

Currently, there are two key types of ANC: Feedforward ANC and Feedback ANC. Feedforward ANC places the microphone on the outside of the earbud or hearing aid to detect external noise (Fig. 2).

This greatly shortens the lag time between the cancellation wave and incoming noise as the system can process and respond to the noise before it enters into the ear canal. However, feedforward ANC is less effective at handling unpredictable or rapidly changing sounds, such as sudden wind gusts, which can exceed its adaptive processing capabilities (Thomas, 2024).

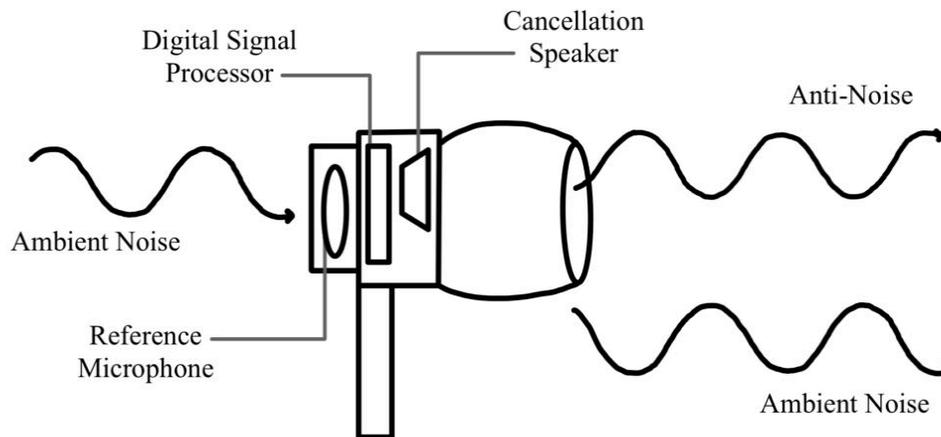
Figure 2
Example of feedforward system



Note. Reprinted from “Engineering Healthy Silence: Using Noise-Cancelling Headphones to Block Harmful Sound” by R. Pak, 2024 (<https://illuminate.usc.edu/engineering-healthy-silence-using-noise-cancelling-headphones-to-block-harmful-sound/>).

Feedback ANC systems position the microphone inside the ear of the device (Fig. 3). This approach is highly effective in canceling out a wide range of noises, as it can directly analyze the sound after it reaches the ear, allowing for more precise cancellation. However, this can introduce greater lag time due to additional processing required (Thomas, 2024).

Figure 3
Example of feedback system



Note. Reprinted from “Engineering Healthy Silence: Using Noise-Cancelling Headphones to Block Harmful Sound” by R. Pak, 2024 (<https://illuminate.usc.edu/engineering-healthy-silence-using-noise-cancelling-headphones-to-block-harmful-sound/>).

Modern-day headphones employ a hybrid of both. Two microphones will be present in the device, with one on the inside and one outside. This dual-microphone system offers the best of both worlds: shortened feedback time and reduced likelihood of ambient sounds breaking through. While hybrid ANC systems perform well in controlled environments, they still struggle with unpredictable sounds. The presence of multiple overlapping noise types can challenge this system's accuracy, potentially resulting in underperformance where unwanted noise persists, or overly aggressive cancellation, which may remove critical sounds such as speech.

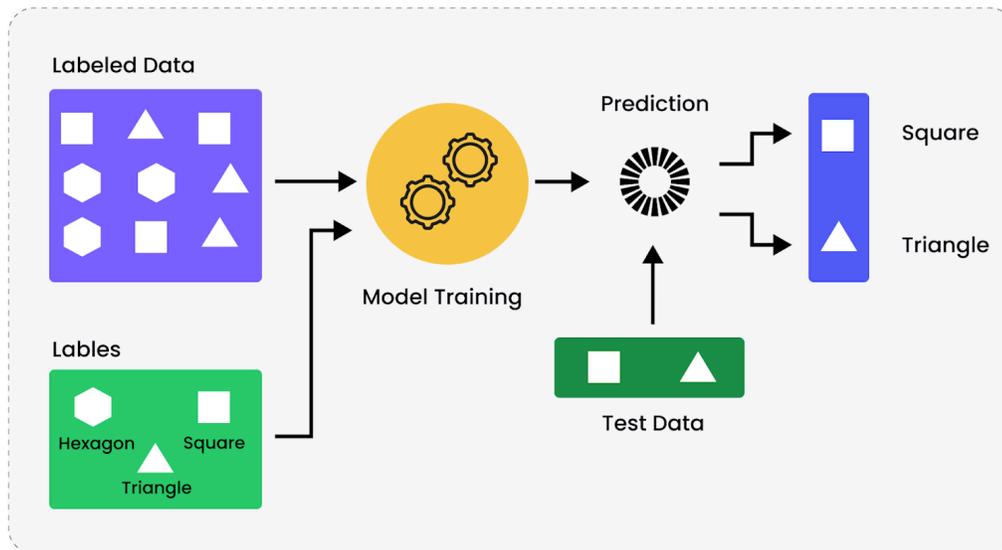
MACHINE LEARNING

Machine learning (ML) is a subset of AI that focuses on developing algorithms and statistical models that enable computers to perform tasks without explicit instructions (Brown, 2021). This sounds complex, but essentially it is the idea of pattern recognition. However, for machines, pattern recognition does not come easy. Hence, as machines are good at crunching huge amounts of numbers, these systems learn patterns and make decisions based on big data (França et al., 2021). At its core, machine learning involves training a model on a dataset to recognize patterns and make predictions or decisions. For hearing aids, this means gathering large quantities of data from the user's environment and listening preferences to optimize sound quality in real-time. Unlike traditional hearing aids, which rely on fixed algorithms, ML-powered hearing aids can continuously improve as more data comes in.

There are three types of machine learning: supervised, unsupervised, and reinforcement learning. Supervised learning involves training a model using a labeled dataset (Fig. 4), enabling it to make

accurate classifications or predictions based on that data (Coursera Staff, 2024). For example, a model is trained with thousands of labeled sets of dog sounds, learning to associate that specific noise pattern with a dog. Over time, the model becomes capable of identifying and classifying similar sounds in real-world environments. This capability is particularly valuable in hearing aids, as it enables the system to differentiate between various background noises and adjust its processing accordingly.

Figure 4
Depiction of model training



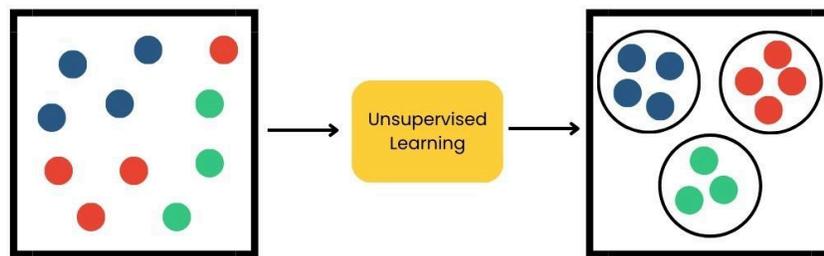
Note. The dataset contains labeled shapes: a square, a triangle, and a hexagon. In the labeled data set are the different shapes. The model depicted will go through the dataset recognizing that the object with 3 sides is a triangle, 4 sides square, and 6 sides hexagon. Testing the model out with a square and a triangle, the model can decipher that the first shape is a square and the second a triangle. Reprinted from “Semi-Supervised Learning Explained: Techniques and Real-World Applications” by R. Panarin, 2024 (<https://maddevs.io/blog/semi-supervised-learning-explained/>).

Unsupervised learning operates without a predefined training set (Fig. 5). Instead, the model analyzes incoming unlabeled data. It will identify patterns on its own and make decisions accordingly (Coursera Staff, 2024). Unsupervised learning can be used for sound environment clustering, or in simple terms, grouping specific sounds to an environment. For example, an AI could collect audio data from various environments- such as a cafe, an office, or a park-and without being explicitly told, distinguish what each setting is. Over time, the model would identify recurring sound patterns and group them into a specific environment. This, in turn, allows the AI to automatically recognize and adapt to

SUMMER 2025

different environments. This makes it particularly effective in real-world settings, where different environments exhibit distinct sound characteristics.

Figure 5
Unsupervised learning model



Note: Three different color circles are depicted. The model has no clue on what it needs to do with all these circles. Eventually, in the final square box, it groups circles by their colors, forming a pattern. Reprinted from “Introduction to Unsupervised Learning” by Rohit M., 2023 (<https://www.bombayssoftwares.com/blog/introduction-to-unsupervised-learning>).

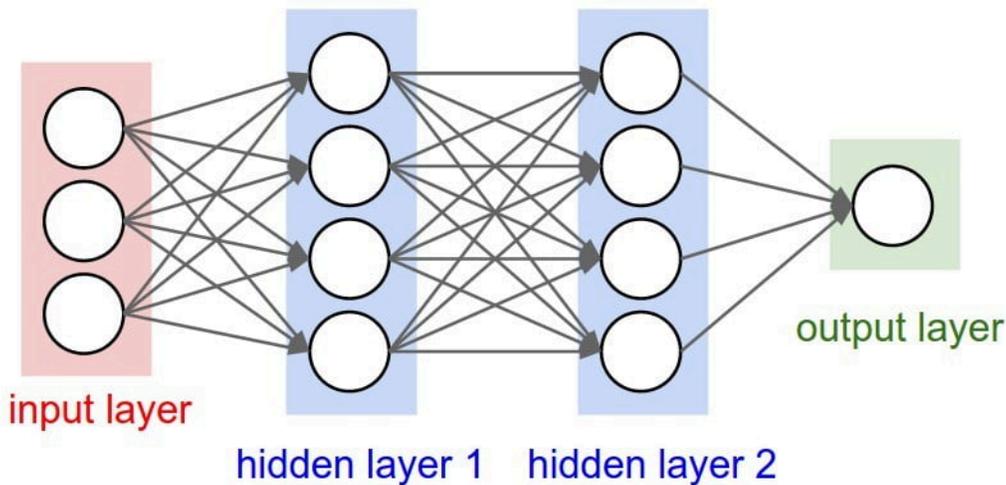
Reinforcement learning uses a reward system. Every time it does something to the user’s liking, the AI is “rewarded” (Coursera Staff, 2024). This is especially helpful over time when an AI can prioritize speech frequencies in noisy environments based on user satisfaction. As machine learning algorithms continue to improve, such as ChatGPT, a hearing aid machine learning algorithm can also improve. This could hypothetically be achieved through the use of user-prompted surveys that will ask the user to rate the clarity of speech in a noisy environment. If the user believes the noise cancellation was enough, they could “reward” the machine learning model with a positive survey so the ML model continues to use that same level in that specific situation.

One of the most notable benefits of machine learning in hearing aids is its ability to enhance noise reduction and speech clarity. Traditional hearing aids often struggle to distinguish between speech and background noise, causing poor performance in noisy environments (Food and Drug Administration, 2022). ML algorithms, however, can analyze sound patterns in real-time, utilizing unsupervised learning, to distinguish different types of noises. For example, Amazon’s Alexa demonstrated that ML-powered speech recognition had only a 6.2% error rate (Raju et al., 2019). This ability of speech recognition will be especially helpful for users in environments such as a noisy restaurant or a mall.

DEEP NEURAL NETWORKS

Deeper into machine learning are deep neural networks, a subset of machine learning (Google, 2025). Deep Neural Networks (DNNs) operate by processing input data through multiple hidden layers, with each layer containing nodes (neurons) connected by weighted links (Hardesty, 2017). These weights indicate the importance of a given feature. As data passes through the layers, each node calculates a weighted sum of its inputs, and if that weight exceeds a defined threshold, that node activates and passes the signal to the next layer. If not, a different node will activate. During training, the DNN network will automatically adjust these weights to improve its predictions (Hardesty, 2017). For example, in a restaurant setting, an audio file may contain various background noises such as forks clinking, water being sipped, waiters writing down orders, and the fire crackling in the kitchen. The DNN will learn to recognize patterns in the audio. It identifies and distinguishes critical sounds (like a person speaking) and background noise. As the network processes the sound, it combines the learned weights to help determine how much noise cancellation should be applied. If the environment is excessively loud, the model might use a stronger level of noise reduction and the vice versa applies. While this is a simplified explanation, it captures the essence of how DNNs enable hearing aids to determine noise suppression levels..

Figure 6
Deep neural network

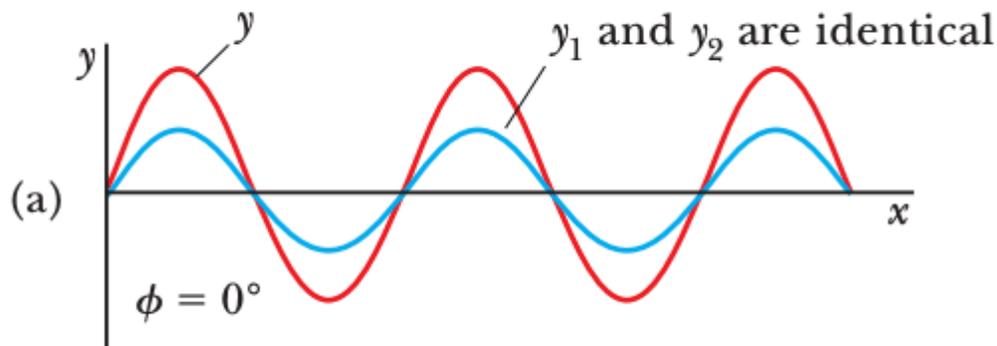


Note. The diagram illustrates an input layer connected to two hidden layers. Each input node contributes to a combination of nodes in the hidden layers, ultimately leading to the output layer. Reprinted from “What’s a Deep Neural Network? Deep Nets Explained” by J. Johnson, 2024 (<https://www.bmc.com/blogs/deep-neural-network/>).

THE LIMITS OF NOISE CANCELLATION

Even with AI, there are inherent limits to the extent of noise cancellation. As discussed in the DNNs section, noise cancellation can operate at varying levels. However, when multiple sounds occur simultaneously, it is not yet feasible to completely eliminate background noise while only preserving only speech. To illustrate simply, imagine two simultaneous sounds: a car rolling by (the louder sound) and rain hitting an umbrella (the quieter one). If the goal is to cancel out the rain sound by using a perfectly inverted signal, that sound can be effectively suppressed. However, this presents a challenge. The anti-sound wave produced does not have a sufficient amplitude to counter the louder car sound (Fig. 7). Consequently, the sound waves do not fully cancel out; while partial noise reduction occurs, leakage will occur (Nave, 2017).

Figure 7
Noise-cancelling anti-sound wave



Note: The blue anti sound wave does not fully match the peak amplitude of the red sound wave.
 Reprinted from "Engineering Healthy Silence: Using Noise-Cancelling Headphones to Block Harmful Sound" by R. Pak, 2024
 (<https://illuminate.usc.edu/engineering-healthy-silence-using-noise-cancelling-headphones-to-block-harmful-sound/>).

Now, consider a different scenario where the goal is to cancel out the car sound. If the anti-noise wave has a greater amplitude than the original sound wave, it will overpower the target signal. This results in a new sound wave of the same frequency with an amplitude equal to the difference between original and anti-noise waves (Science Learning Hub, 2019). Granted, the resulting sound will still be audible, but it will be quieter than the original sound (Science Learning Hub, 2019). Either way there will be a leakage of sound which poses a problem. This is the reason why a high amplitude anti-wave cannot be used to resolve all problems. In the process of canceling out the louder noise, the system could

SUMMER 2025

inadvertently cancel out the useful speech, as the anti-wave noise might completely overpower the critical sound wave.

MACHINE LEARNING AND THE ANTI-SOUND WAVE

This section presents a simplified example of how a machine learning model could be applied to generate an anti-wave to minimize background noise while still maintaining speech clarity. Although simplified, the example is intended solely to act as only a hypothetical illustration of the process. The model described aims to predict the optimal anti-noise wave that minimizes interference while preserving speech quality.

To establish a few ground rules for this hypothetical model, a few simplifications will be made. Instead of using raw decibel numbers, a scale of 1-10 with 1 being the weakest and 10 being the greatest will be utilized. Hearing difficulty level will be the only exception to the 1-10 scale level; it will be classified on a scale of 1-5. For simplicity, only two types of noise will be considered instead of five. Additionally, speech clarity will be measured on a scale from 1 to 10.

The target output of this model is the predicted anti-noise wave level, which will also be categorized on a scale from 1 to 10. In this hypothetical model, an anti-wave of 1 would cancel out completely a noise at scale 1. The objective of the anti-wave is to keep the disturbance level, a simplified hypothetical variable, minimized. The disturbance level would be calculated by taking the difference between the anti-wave level and the noise level. For example, if the anti-wave level is 6 and the noise level is 5, the disturbance level would be 1. Similarly, if the anti-wave level is 5 and the noise level is 6, the disturbance level would also be 1.

Table 1
Example ML dataset

	Speech Level	Noise Level 1	Noise level 2	Hearing Difficulty	Anti-Wave Level
1	Trial 7	2	2	4	2
2	Trial 4	3	4	1	3
3	Trial 7	5	5	5	2
4	Trial 9	7	2	3	4

SUMMER 2025

Through the example dataset (Table 1), the model would identify a consistent pattern indicating that the speech level should never be above the anti-wave level. As for each of the data set values, the target anti-wave value is never greater than the speech level (e.g. Trial 1: $7 > 2$, Trial 2: $4 > 3$, Trial 3: $7 > 2$, and Trial 4: $9 > 5$). Consequently, the model learns that the chosen anti-noise level cannot be equal to or exceed the speech level.

Subsequently the model identifies a correlation between the dataset's hearing difficulty level and anti-wave level. Suppose the model is given a high hearing difficulty level of 5. In this case, the anti-wave level should be significantly lower than the speech level. By analyzing the example dataset (Table 1), the model is able to infer the underlying pattern. It observes how in trials 1, 2, 3, and 4, the anti-wave level is less than or equal to the difference of the speech level. For example, In trial 3, the speech level is 7 and the hearing difficulty level is 5, resulting in an anti-wave level of 2. Based on this pattern, the model formulates the idea that the anti-wave level must be less than or equal to the speech level minus the hearing difficulty level.

Next, the model shifts its attention to the noise variables. Sifting through Trial 1, it simply determines that an anti-wave of level 2 works. It meets all prior requirements shown earlier, and it realizes that, as $2 - 2 = 0$ for both. It comes to the conclusion that the disturbance level is 0. From this trial, the model infers that the anti-noise wave level closely corresponds to the noise level. Moving on to Trial 4, it sees a noise level of 7, a noise level of 2, and an anti-noise level of 4. From this trial, the model identifies a pattern of minimizing disturbance: it calculates the average of the two noise values by summing them, dividing by two, and rounding down. Albeit the caveat is that it still has to abide by the previous patterns it recognized. This allows the model to now minimize the noise and finally produce the optimal anti-noise wave.

This all culminates into a finished model that, given a new trial as a test, can predict the output on its own. Take, for example, in an environment where the speech level is 10, 6 and 4 for the noise levels, and 5 for the hearing difficulty level. The Model will be able to discern that in this environment, an anti-wave level of 5 should be the effective output. The model recognizes that the anti-wave level cannot exceed 10. Given a hearing difficulty level of 5 and a speech level of 10, it determines that the difference between these values allows for an anti-wave range of 0 to 5. To select the optimal anti-wave level, the model minimizes disturbance by averaging the two noise levels—6 and 4—resulting in an anti-wave value of 5.

It needs to be reiterated that this example only illustrates a highly simplified version of how a machine learning model could be used to predict the anti-wave level required to minimize disturbance while maintaining speech clarity. This is not an actual model rather a hypothetical representation of one. It is also important to note that in a real-world scenario, the model would need to account for a much greater number of variables. It might have to operate using unsupervised learning techniques, particularly in cases where raw, unlabeled data is used. An example would be audio files where the ML would have to go a step further and distinguish between noise and speech patterns. The complexities of

SUMMER 2025

sound and machine learning algorithms make real-world implementations far more difficult to accomplish.

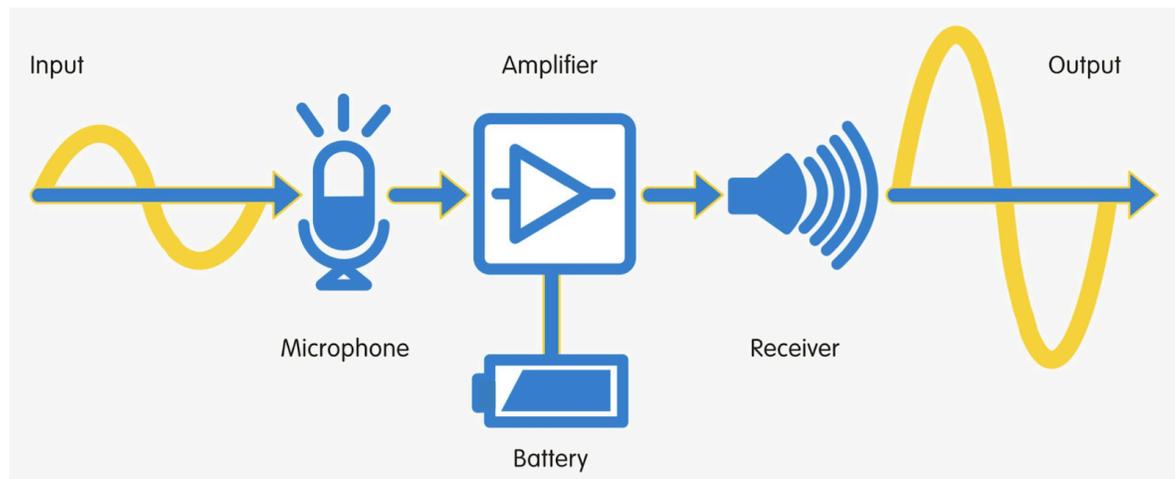
CURRENT COMPANIES

As the development of AI in hearing aids is relatively new, few companies have entered the space. The most prominent being Starkey, utilizing their Edge AI. Edge AI uses a DNN to distinguish “non-stationary source noise, and in turn better reduce noise levels for these noise types” (Betlehem et al., 2024). Still the challenge with multiple noises coming in at once remains; however, Edge AI has come up with a better way to manage the scenario. It can distinguish between two types of noise: stationary and nonstationary. With stationary noise being, for example, a constant hum of a refrigerator, and non-stationary noise being the clanging of pots and pans. Leveraging this capability, Edge AI is also able to “give more relief in environments where a mix of stationary and non-stationary noise is present” (Betlehem et al., 2024). According to Starkey, their DNN technology is 30% better at identifying speech than the legacy approach, traditional way, to distinguish between speech and noise. This improvement in noise reduction is evident from the testing done at Starkey. In a benchtop evaluation comparing “Road Noise with Bumps” and “Kitchen Fan with Clanking Dishes”, the DNN-based noise reduction implemented in Edge AI achieved reductions of up to 2 dBA under default settings and up to 3.5 dBA with enhanced settings (Betlehem et al., 2024) Even in more subtle noise environments, such as the “Kitchen Fan with Clanking Dishes” scenario, DNN still outperformed traditional methods with a 1.5 dBA improvement. This demonstrates the idea that over a variety of complex noise scenarios, Starkey's overall noise reduction significantly surpassed the traditional method (Betlehem et al., 2024). According to Starkey, a perceptual study conducted revealed that participants preferred the DNN based hearing aid over the traditional version. In a study involving 14 experienced hearing aid users, the vast majority of participants preferred the DNN-based algorithm over the legacy approach in terms of perceived noise reduction, especially in challenging noise environments such as restaurants, bars, and other similar social settings (Betlehem et al., 24). The strongest preference was observed most in challenging speech-in-noise environments, particularly when background noise had a signal-to-noise ratio (SNR) of -5 dB. This, in basic terminology, means that the background noise was 5 decibels higher than the actual speech. While individual preferences varied, some participants favored different options in speech-in-noise conditions despite having similar hearing loss profiles. Though, the overall trend indicated a clear advantage for Edge AI's DNN-based technology.

EXTRA BENEFITS

Moving beyond noise reduction, amplification remains to be a fundamental function of hearing aids. A key feature frequently highlighted by Starkey's Edge AI is its ability to differentiate speech from background noise. The primary purpose of a hearing aid is to amplify speech. Hearing aids traditionally enhance sound by a three-step process (Fig. 8) where a microphone takes in sound, converting it into a digital signal, using an amplifier to increase the digital sound, then relaying it back to the user (John Hopkins, n.d.).

Figure 8
Sound wave amplifier



Note: A sound wave is shown passing through a microphone that picks up the sound. It then is passed through an amplifier. This amplifier subsequently relays the sound wave to a receiver at a higher amplitude, seen with the greater magnitude sound waves (higher amplitude, louder the sound). Reprinted from “How do hearing aids work?” T. Sieber, 2023 (<https://www.soundguys.com/how-do-hearing-aids-work-57482/>).

The primary limitation of this system is that it amplifies all sounds indiscriminately, making it difficult to isolate the desired audio in environments with multiple overlapping sounds. However, if an AI can distinguish speech, it is plausible that it can selectively increase the sound of the speech while minimizing noise. Although this remains a hypothetical application of AI in hearing aids, its potential merits serious consideration. Starkey AI briefly explained this concept of “improved speech presence prediction” (Bethlehem et al., 2024).

Potential for reduced cost is another worthy benefit. As noted in the introduction, cost is a major barrier to the average American. With the average price for a quality hearing aid being around \$6000 in 2016, it creates a major deterrent for those seeking help (Bluestein et al., 2016). Given inflation and other economic factors, this cost is likely even higher today. In addition to the upfront cost, the cost to adjust and adapt hearing aids to the user's preferences is not negligible either. All degrees of hearing loss are not the same. Some may need the sound of everything to be adjusted ever so louder than others (Traynor et al., 2002). Many individuals have their own personal “style” of hearing preferences (Traynor et al., 2002), which often requires them to visit an audiologist for hearing aid adjustments. Since these adjustments typically occur every 6 to 12 months, the associated costs can accumulate rapidly (Susan Rogan Hearing, 2022). Average cost of an audiologist visit in New Jersey is in the range of \$96 to \$142 a

SUMMER 2025

visit (Sidecar, n.d.). With AI being capable of adjusting tones and volume level, this hypothetical “personal style” can be achieved. This theoretically eliminates the need to visit an audiologist either completely or, more realistically, reduces visits to once every two years.

CONCLUSION

There is no doubt that the potential of AI in hearing aids can revolutionize the hearing aid industry. By using AI, more specifically machine learning, to its full potential, there will be substantial advancement in hearing aid technology. By focusing on personalized sound profiles, dynamic noise management, and real-time adaptability, AI can greatly improve speech clarity in different sound environments. This, in turn, will enable AI to have the capacity to improve quality of noise cancellation, enhance speech clarity, and even lower hearing-related costs. These key advancements come at a crucial time. As previously noted, hearing loss is an increasingly prevalent issue, with 10% of Generation Z and 17% of Millennials already experiencing some form of hearing impairment. This is a significantly higher rate than that of older generations at the same age (American Hearing Audiology, 2023). Whether this rise is attributed to the use of personal audio devices or other environmental factors, it is clear that hearing challenges will be a pressing issue for future generations that must be addressed. While the idea of a world without hearing loss remains an optimistic future, advancements in technology can continue bridging the gap in the meantime. The tools and technologies we create one day may continue to improve the quality of life for those affected by hearing challenges. It is up to us to address the current imperfections, and through innovation, we can and will eliminate this issue.

REFERENCES

American Hearing Audiology. (2023, June 17). *Hearing loss among millennials and Gen Z*. American Hearing Audiology. <https://americanhearing.us/hearing-loss-among-millennials-gen-z/?srsltid=AfmBOodlbrGl-7gKIGWAVY-30ERYL7I0XJnrUNOBgqRrDD26WphxGcT>

Betlehem, T., Mishra, P., Xu, J., Marquardt, D., & McKinney, M. (2024). Leveraging DNN in Edge AI 1. In *Starkey*. <https://cdn.mediavalet.com/usil/starkeyhearingtech/DPIrh4nT5UWv37iI1gwcqA/5FtECxMyD0qjVCcNhAH4Sw/Original/DNN%20Assisted%20MNR%20Whitepaper.pdf>

Blustein, J., & Weinstein, B. E. (2016). Opening the market for lower cost hearing aids: Regulatory change can improve the health of older Americans. *American Journal of Public Health*, 106(6), 1032–1035. <https://doi.org/10.2105/AJPH.2016.303176>

Brown, S. (2021, April 21). *Machine learning, explained*. MIT Sloan; MIT Sloan School of Management. <https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained>

SUMMER 2025

Coursera Staff. (2024, April 3). *3 types of machine learning you should know*. Coursera. <https://www.coursera.org/articles/types-of-machine-learning#>

França, R. P., Borges Monteiro, A. C., Arthur, R., & Iano, Y. (2021, January 29). *Chapter 3 - An overview of deep learning in big data, image, and signal processing in the modern digital age* (V. Piuri, S. Raj, A. Genovese, & R. Srivastava, Eds.). ScienceDirect; Academic Press. <https://www.sciencedirect.com/science/article/abs/pii/B9780128222263000039>

Google. (n.d.). *Deep learning vs machine learning*. Google Cloud. <https://cloud.google.com/discover/deep-learning-vs-machine-learning?hl=en>

Hardesty, L. (2017, April 14). *Explained: Neural networks*. MIT News; Massachusetts Institute of Technology. <https://news.mit.edu/2017/explained-neural-networks-deep-learning-0414>

John Hopkins Medicine. (2020). *Types of hearing loss*. John Hopkins Medicine. <https://www.hopkinsmedicine.org/health/conditions-and-diseases/hearing-loss/types-of-hearing-loss>

Johns Hopkins Medicine. (n.d.). *How do hearing aids work*. Www.hopkinsmedicine.org. <https://www.hopkinsmedicine.org/health/conditions-and-diseases/hearing-loss/how-do-hearing-aids-work>

Johnson, J. (2020, July 27). *What's a deep neural network? Deep Nets Explained*. BMC Blogs. <https://www.bmc.com/blogs/deep-neural-network/>

Mroz, M. (2022, June 28). *Hearing aid prices - How much do hearing aids cost?* Healthy Hearing. <https://www.healthyhearing.com/help/hearing-aids/prices>

National Institute on Deafness and Other Communication Disorders. (2024, September 20). *Quick Statistics About Hearing*. NIDCD. <https://www.nidcd.nih.gov/health/statistics/quick-statistics-hearing>

Nave, R. (2017). *Interference of Sound*. Gsu.edu. <http://hyperphysics.phy-astr.gsu.edu/hbase/Sound/interf.html>

Pak, R. (2024, April 1). *Engineering healthy silence: Using noise-cancelling headphones to block harmful sound - USC Viterbi School of Engineering*. Illumin.usc.edu. <https://illumin.usc.edu/engineering-healthy-silence-using-noise-cancelling-headphones-to-block-harmful-sound/>

Panarin, R. (2024, August 8). *What semi-supervised learning is and how does it work*. Custom Software Development Company; Mad Devs Group LTD. <https://maddevs.io/blog/semi-supervised-learning-explained/>

Raju, A., Filimonov, D., Tiwari, G., Lan, G., & Rastrow, A. (2019, July 2). *Scalable multi corpora neural language models for ASR*. <https://arxiv.org/pdf/1907.01677>

Rohit, M. (2023, September 9). *Introduction to unsupervised learning: Overview and applications*. Bombaysoftwares.com. <https://www.bombaysoftwares.com/blog/introduction-to-unsupervised-learning>

Science Learning Hub. (2019, September 13). *Sound – wave interference*. Science Learning Hub. <https://www.sciencelearn.org.nz/resources/2816-sound-wave-interference>

Sidecar. (n.d.). *Cost of audiologist visit by state | Sidecar Health*. Cost.sidecarhealth.com. <https://cost.sidecarhealth.com/c/audiologist-visit-cost>

Sieber, T. (2023, March 7). *How do hearing aids work?* SoundGuys. <https://www.soundguys.com/how-do-hearing-aids-work-57482/>

Susan Rogan Hearing. (2022, June 15). *3 FAQs about hearing aid fittings*. Susan Rogan Hearing. <https://susanroghanhearing.com/3-faqs-about-hearing-aid-fittings/>

Sygrove, C. (2024, April 25). *Hearing aids with artificial intelligence (AI): Review of features, capabilities and models that use AI and machine learning*. Hearing Tracker. <https://www.hearingtracker.com/resources/ai-in-hearing-aids-a-review-of-brands-and-models>

Thomas, C. (2024a, July 22). *Active noise-canceling (ANC) technology types explained*. SoundGuys. <https://www.soundguys.com/noise-canceling-anc-explained-28344/>

Thomas, C. (2024b, July 22). *How do noise-cancelling headphones work?* SoundGuys. <https://www.soundguys.com/how-noise-cancelling-headphones-work-12380/>

Traynor, R. M., & Holmes, A. E. (2002). Personal style and hearing aid fitting. *Trends in Amplification*, 6(1), 1–31. <https://doi.org/10.1177/108471380200600102>

Victory, J. (2024, February 26). *Hearing loss statistics at a glance*. Healthy Hearing; Healthy Hearing. <https://www.healthyhearing.com/report/52814-Hearing-loss-statistics-at-a-glance>

World Health Organization. (2025, February 26). *Deafness and hearing loss*. WHO; World Health Organization: WHO. <https://www.who.int/news-room/fact-sheets/detail/deafness-and-hearing-loss>