

Brain-Controlled Hearing Aids Could Cut through Crowd Noise

A prototype detects whom you are listening to and amplifies only that speaker's voice; a potential solution to the "cocktail party problem"

By Lydia Denworth on May 15, 2019

At a crowded party or a noisy restaurant, most of us do something that is remarkable. Out of all the voices surrounding us, our brains pick out the one we want to hear and focus on what that person has to say. People with hearing loss are not so fortunate. Noisy situations are especially difficult for them and hearing aids and cochlear implants do not help much. Such technology generally either amplifies all voices or mashes them together so they are indistinguishable. The question of how the brain manages the trick of hearing in noise is known as the "cocktail party problem." It is a puzzle that has bedeviled auditory scientists for decades and limited the solutions they have to offer. But researchers have just taken a major step forward toward helping people hear in noise. In a paper published on May 15 in *Science Advances*, engineers from Columbia University's Zuckerman Institute revealed an experimental technology that could lead to a brain-controlled hearing aid. Their proof-of-concept device uses artificial intelligence to separate voices and compare them with a listener's brainwaves to identify and amplify the speaker to whom that listener is paying closest attention. Nima Mesgarani of Columbia University's Zuckerman Institute, the senior author on the paper, has been working on aspects of the same problem since 2012 when he first discovered it was possible to figure out which voice a listener was focused on by monitoring brainwaves.

In 2017, he developed technology that could pull one voice from many, but only if the system was trained to recognize that particular speaker—a severe limitation in realworld communication. Now Mesgarani and his colleagues have achieved a significant step forward by using brainwaves to decode whom you are listening to and then separating the interlocutor's voice without the need for training. "To remove that barrier," he says, "is a pretty big breakthrough."

"It's a beautiful piece of work," says auditory neuroscientist Barbara Shinn-Cunningham, director of the Neuroscience Institute at Carnegie-Mellon University, who was not involved in the research. Auditory neuroscientist Andrew Oxenham of the University of Minnesota, who has studied the cocktail party problem for years, says, "This brings the whole field closer to a practical application, but it's not there yet."

What Mesgarani and his colleagues have created is an algorithm, and they have tested it only in epilepsy patients undergoing brain surgery. Such patients provide a rare opportunity for scientists to put electrodes directly into human brains. From a loudspeaker in front of the participants, Mesgarani and his colleagues played two

voices (one male, one female) speaking simultaneously. They instructed participants to focus first on one and then the other. The Columbia engineers fed the sound of the voices and the electrical signals from the patients' brains into their algorithm, which sorted the sounds, amplified the attended voice and attenuated the other. "These two inputs go inside this box, and what comes out of it is the modified audio in which the target speaker is louder," Mesgarani says.

Although using brainwaves to follow auditory attention is an impressive achievement, the real advance has to do with the algorithm. It uses a sophisticated form of artificial intelligence known as a deep attractor network to separate unknown speakers automatically and in real time. Such neural network models, developed within the last four years, look for statistical regularities in increasingly complex layers of computations to determine which parts of a sound mixture belong together. "Deep learning is the secret sauce that made [this] possible," Mesgarani says.

Credit: Columbia University's Zuckerman Institute

It doesn't matter that neuroscientists still haven't fully worked out how the brain hears in noise. "We are not trying to simulate the brain," Mesgarani says. "We are just trying to solve the cocktail party problem." They trained the algorithm with far more examples of human speech than any person would hear in a lifetime. Then they gave it the task of analyzing the detailed, often overlapping information in the spectrograms, or acoustic signatures, created by multiple speakers' voices and separating them into distinct streams of sound. Graphically represented, the paper shows two combined voices as a haze of red and blue dots. Once separated, one voice is a cluster of red dots, the other blue. There is still an element of mystery in how exactly the algorithm does this. "Our guess is that it uses the spectral and temporal information, common onsets and offsets [speech characteristics], and harmonic structures," Mesgarani says. "We tell it that this cloud of red and blue should become separable. It figures out somehow magically this transformation, and suddenly you have two clouds."

Credit: Nima Mesgarani Columbia University's Zuckerman Institute

Considerable challenges remain before this technology can be used in an actual hearing aid. Mesgarani estimates it will be at least another five years. Of course, a marketable device requires a noninvasive technique for generating EEG recordings of brainwaves. Several scientists, including Mesgarani, have shown that in-the-ear or around-the-ear hearing aids with electrodes can work, although they generate a far less precise signal. And while powerful, the algorithm is still not yet successful 100 percent of the time.

In all probability, the first devices to use this technology will help people with mild to moderate hearing loss. "You probably need some residual hearing," Mesgarani says. "As long as you can track the ups and downs of [one] voice, that

would be the kind of signature that this technology would look for. The talker separation algorithm alone could prove helpful without monitoring brainwaves at all, says electrical engineer Mario Svirsky of New York University's Langone Medical Center. "I envision a smartphone app that talks to your hearing aid," he says. "The app shows you icons for different talkers. If you click an icon, then that talker is preferentially amplified and the others attenuated." As for a true brain-controlled hearing aid, Svirsky fears that the costs may outweigh the benefit and is skeptical one will ever be implemented. But he remains enthusiastic about Mesgarani's work. "The whole idea of having a mind-reading hearing aid is fascinating," Svirsky says. "It's not just science fiction. This research has shown that it is at least a plausible possibility."

Lydia Denworth

Lydia Denworth is a Brooklyn-based science writer and author of *I Can Hear You Whisper: An Intimate Journey through the Science of Sound and Language* (Dutton, 2014). She is working on a book about the science of friendship.

Credit: Nick Higgins

Recent Articles

Children Change Their Parents' Minds about Climate Change

"Hyperscans" Show How Brains Sync as People Interact

Forestalling a Fatal Decision

<https://www.scientificamerican.com/article/brain-controlled-hearing-aids-could-cut-through-crowd-noise/?redirect=1>