

**NEW BRAIN-CODING TECHNOLOGY** brings us one step closer to mind reading.

# WE WANT TO HEAR YOUR THOUGHTS

BY AVERY HURT

**It didn't take long for Edward Chang** to see the implications of what he was doing. The neuroscientist and brain surgeon at the University of California, San Francisco, was studying the brain activity behind speech, that precise and delicate neural choreography by which lips, jaw, tongue, and larynx produce meaningful sounds.

By implanting an array of electrodes between the outer and inner membranes of the brain, directly over the area of the brain that controls speech, he and his team were able to detect distinct patterns of brain activity associated with specific sounds, each vowel and consonant, each *duh*, *guh*, *ee*, and *ay* sound that combine to form words.

"We realized that we had a code for every speech sound in the English language," Chang says. And that realization opened up some astonishing possibilities.

In a series of papers published between 2019 and 2021, Chang and his team demonstrated how they could use machine learning, a form of artificial intelligence, to analyze the patterns. They immediately saw the potential benefits for people who've lost the ability to speak because of brain-stem stroke, cerebral palsy, ALS, or other forms of paralysis: Once people's words and sentences are reconstructed through analysis of those brain patterns, the words can be displayed as text on a screen. More recently, the researchers demonstrated that the words a person is trying to say can even be translated into a computer-generated voice and facial movements on an on-screen avatar, enabling a paralyzed person to communicate not just with speech, but with facial expressions as well.

Meanwhile, researchers at the University of Texas at Austin are working on a less invasive method for

peering into the mind. A team led by Alexander Huth, a computational neuroscientist, uses functional magnetic resonance imaging (fMRI), rather than implants, to monitor brain activity. Then, much like Chang's group, they use a machine-learning system called a "semantic decoder" to match each word or phrase with a particular pattern of brain activation.

"Basically, we build a model of a person's brain. And then when we get new brain recordings from that person, we can use the model to generate a sequence of words that predicts what the user is hearing or imagining," explains Jerry Tang, a neuroscientist in Huth's lab and lead author on many of the studies on this technology. "It's not like some of the other studies that look at the words they're attempting to say. It's actually their thoughts, what they're imagining."

A paper by Tang and colleagues published in *Nature Neuroscience* in May 2023 gave an example. When one participant listened to the words, "I didn't know whether to scream, cry, or run away. Instead, I said, 'Leave me alone!'", the AI decoded the thought as: "Started to scream and cry, and then she just said, 'I told you to leave me alone. You can't hurt me.'"

"It's not perfect, but it's shockingly good for using fMRI," Huth said at a February 2024 meeting of the National Institutes of Health's Neuroethics Working Group, where he discussed his and his team's work.

**USING AN FMRI** scanner, researchers (from left to right) Alex Huth, Shailee Jain, and Jerry Tang gathered the brain activity data used to train their semantic decoder (below).



*Shocking* is the right word. Huth told *Science* magazine that on seeing that this actually worked, his first thought was, “Oh my God, this is kind of terrifying.”

**TERRIFYING** or not, all of this research to date has not crossed the threshold into mind reading — at least not yet. The researchers are careful to point out that this method was intended to work only with cooperative participants. The volunteers in Tang’s studies spent 16 hours having their brains scanned while they listened to stories from the podcasts *Moth Radio Hour* and *Modern Love*. This provided researchers with an abundance of data about the activity of the volunteers’ brains while listening to spoken words. The AI then drew on this rich database to seek



patterns it could match to specific words and sentences. But the data using this method did not seem to be transferable. In other words, the AI could not decode the thoughts of another person based on training data from a participant’s brain.

Nonetheless, researchers are aware of the implications. “We thought pretty deeply about what this could mean,” Tang says. Ultimately, they established that people could consciously control which of their thoughts are decoded. For example, Tang explains that if someone is hearing two different stories at the same time, they can control which of the stories is decoded. “We took this to suggest that you can only decode what [a person] is actively thinking about.”

Still, what worries some people isn’t what the technology can do now, but what it may one day be able to do.

Tang acknowledges that this technology is in its infancy, and it’s conceivable that it could get to the point where it could be used to read a person’s thoughts against their will. “There’s a lot more information in these brain scans than we previously thought,” Tang says. “And we don’t know what the ceiling on that is. So, we definitely don’t want to give anyone a false sense of security.”

And indeed, advances in the technology are coming fast — faster than even

its developers expected. In his presentation to the Neuroethics Working Group, Huth shared some as yet unpublished — and surprising — results indicating that the brain data may be transferable after all. The team’s studies have demonstrated that they can, in fact, decode a second person’s thoughts by using the larger dataset from the first person, though the technique requires at least a small amount of training data from the second person. “It’s remarkable what we can do with almost no training data from a person,” Huth said at the meeting. While the accuracy of this method so far is limited, it will improve as computing power increases. “We haven’t hit the limits yet,” he said.

“I think the lesson of the modern era is that generative AI is making everything, including brain decoding, go much faster than people thought,” says Nita Farahany, professor of law and philosophy at Duke University and author of *The Battle for Your Brain: Defending the Right to Think Freely in the Age of Neurotechnology*.

**THE POTENTIAL** privacy implications of developing any form of “thought-reading” technology are easy to imagine. When the technology is limited to medical applications, such

as Chang’s device that helps paralyzed people regain the ability to communicate, those applications are covered by the Health Insurance Portability and Accountability Act (HIPAA). HIPAA imposes strict privacy regulations on health-care providers, and severe penalties for violating them. Researchers, meanwhile, are subject to restrictions imposed by independent committees, called institutional review boards, that must review and approve the ethics of proposed research before it can go forward.

But this technology isn’t solely restricted to neuroscience labs, and the applications aren’t limited to health care. Companies providing this technology for use outside health-care settings aren’t bound by those same privacy regulations or ethical standards.

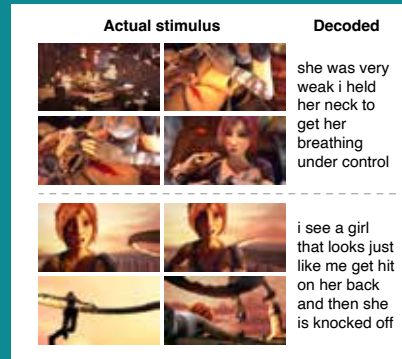
Wearable devices as well as apps that can read and record brain data are already available commercially. Often referred to as “personal neurotech,” these products let you use standalone devices or even your smartphone to record or monitor your own brainwaves in real time to improve your meditation, reduce stress, or enhance focus. And more are on the way. This summer, Apple secured a patent for earbuds that would contain sensors that use electroencephalography (or EEG) to measure and show users the electrical activity in their brain.

Of course, apps and personal tech lag well behind the advances in research labs — none of the commercial devices can predict your thoughts. But Farahany points out that even today, when the commercial technology is fairly primitive, these devices can reveal more about what you’re thinking and feeling than you might be comfortable with.

And the information these devices collect doesn’t belong to you. Remember the small print on consumer agreements that nobody ever reads before clicking “I agree” and downloading the software? Rafael Yuste has read the small print — and it made for more interesting reading than you might expect. Yuste is chair and co-founder of the Neurorights Foundation, an organization dedicated to ensuring the ethical development of neurotechnology. Without exception,

## A DECODER RINGS TRUE

**The semantic decoder** developed at University of Texas at Austin was trained on brain data from volunteers who watched clips from the movie *Sintel* without sound; the decoder captured the gist of the scenes. Volunteers also listened to four audio stories. The color-coded text (bottom) shows words the decoder reproduced exactly, words for which it captured the gist, and also a few errors.



### Actual stimulus

i got up from the **air mattress** and **pressed my face against the glass of the bedroom window** expecting to see eyes **staring back at me but instead finding only darkness**

i didn't know whether to **scream cry** or **run away** instead i **said leave me alone i don't need your help adam disappeared** and i **cleaned up alone crying**

that night i went upstairs to what had been our **bedroom** and **not knowing what else to do i turned out the lights** and **lay down on the floor**

i don't have my **driver's license yet** and i just **jumped out right when i needed to and she says well why don't you come back to my house and i'll give you a ride i say ok**

### Decoded stimulus

i just continued to **walk up to the window** and **open the glass i stood on my toes and peered out i didn't see anything and looked up again i saw nothing**

started to **scream and cry** and then **she just said i told you to leave me alone you can't hurt me anymore i'm sorry and then he stormed off i thought he had left i started to cry**

we got back to my **dorm room** i had **no idea where my bed was** i just assumed i would sleep on it but instead i **lay down on the floor**

she is not ready **she has not even started to learn to drive yet** i had to **push her out of the car** i said **we will take her home now and she agreed**

Exact

Gist

Error



Yuste says, every single one of the agreements for commercial neurotech devices and apps give ownership of all the user's brain data to the company that collects it.

What's more, when agreeing to most of these contracts — and yes, they are contracts — the user allows the company to sell the data to a third party, which is not bound to any agreement the user may have made with the original provider. “In other words,” says Yuste, “the status of brain data in commercial technology could not be less protected.”

**FOR TANG**, the risks of this technology may lie not in underestimating what these devices can do, but rather in overestimating their abilities. He points to polygraph tests, which are generally agreed to be inaccurate in that they appear to measure anxiety rather than deception. The result is a lot of false positives that can, and have, led to the miscarriage of justice. “In the same way, I think it's important that the capabilities of this technology are not overexaggerated,” he says. “It's important to be transparent about exactly what we can do, but also what we can't do, in order to make sure that decoding technology is not misused.”

Though the risks of misuse of this technology are real, the potential benefits are enormous. Chang and Tang are using their discoveries to develop technologies that can help restore the ability to communicate. Marcel Just has an even grander vision: solving the mystery of the human mind. Just, a cognitive neuroscientist at Carnegie Mellon University, helped pioneer the use of fMRI and machine learning to understand how the brain stores and processes concepts and meaning. He likens this technology to the insightfulness provided by the first microscope or first telescope. “It's a scientist's dream,” he says. “It has opened the door to



understanding the nature of human thought.”

Just sees this technology as an aid to building better brains and developing more effective ways to teach and to learn, in the same way advances in exercise physiology have helped develop safer and more effective ways to build better bodies. “We can not only make better athletes, which we've proven we can do, we can make better thinkers,” says Just. “If you know how the brain handles everything in the world, you could teach everything in the world accordingly. You could make educational systems enormously more streamlined.”

Some of the benefits of brain-decoding technology are more practical. For example, wearable sensors that monitor the brains of long-haul truck drivers and alert them when they're too sleepy to drive could be lifesaving — not just for the truckers, but for everyone on the road. In this case, says Farahany, the interest in protecting society from a sleep-deprived driver barreling down the highway likely outweighs truck drivers' concerns about the privacy of their fatigue levels.

Gabriel Lázaro-Muñoz,

a neuroscientist with a background in law and philosophy, works at the Center for Bioethics at Harvard Medical School. He says that finding a solution that balances the need for privacy with the opportunity for medical advances will require that both the public and policymakers become educated about the risks and benefits of this technology.

Yuste and the Neurorights Foundation are working with governments worldwide to include neural rights protections in their constitutions. In 2021,

Chile became the first country to enshrine the right to neural privacy in its constitution. In 2023, the Brazilian state of Rio Grande do Sul incorporated neural rights into its constitution, and the same year Mexico added neural rights to its Charter of Digital Rights. In the U.S., Colorado became the first state to protect neural rights, when, in August 2024, it enacted a law adding biological data — including neural data — to the state's existing privacy protections.

Farahany agrees that protections are needed, but current efforts don't do enough. The neuro rights approach is both overspecified and underinclusive, she says. By focusing exclusively on legislation and constitutional amendments to protect neural data, it neglects to address the ecosystem in which this data is used, such as targeted ads and manipulative technologies.

“We need to have a national conversation about how we want to manage these technologies,” Lázaro-Muñoz says. He seems confident that conversation will happen. “The public generally likes talking about mind-reading technologies.”

“Generative AI is making everything, including brain decoding, go much faster than people thought.” — Nita Farahany

Avery Hurt is a regular Discover contributor.