The New York Times

"When Your Inner Thoughts Are No Longer Just Yours"

By Trevor Paglen
November 12, 2026

Dr. Paglen is an artist and the author of the book "How to See Like a Machine: Images After A.I."

The possibility we could one day read one another's minds is more real than ever today, thanks to rapid advancements in artificial intelligence. The technology is moving fast, and the paranoia that science may open the inner recesses of our minds to manipulation no longer seems so outlandish.

Many researchers have the laudable intention of learning about the brain, but in doing this work they may be opening the door to our inner worlds' being exploited for commercial gain. It's a terrifying future we aren't prepared for.

I learned this for myself a year ago at Cornell University, when I spent three hours, over two days, strapped in an fMRI machine (which can show the real-time changes in blood flow in the brain). Lying in the scanner, I concentrated on a sequence of seemingly random images projected onto a screen a few feet from my eyes, while the machine rancorously probed the gray matter of my brain. Pictures of traffic lights, cars, airplanes and trees were meant to elicit a "boring" response from my amygdala (which governs fight-or-flight responses). Images of grotesque A.I.-generated faces or people making threatening gestures were meant to excite that part of my brain.



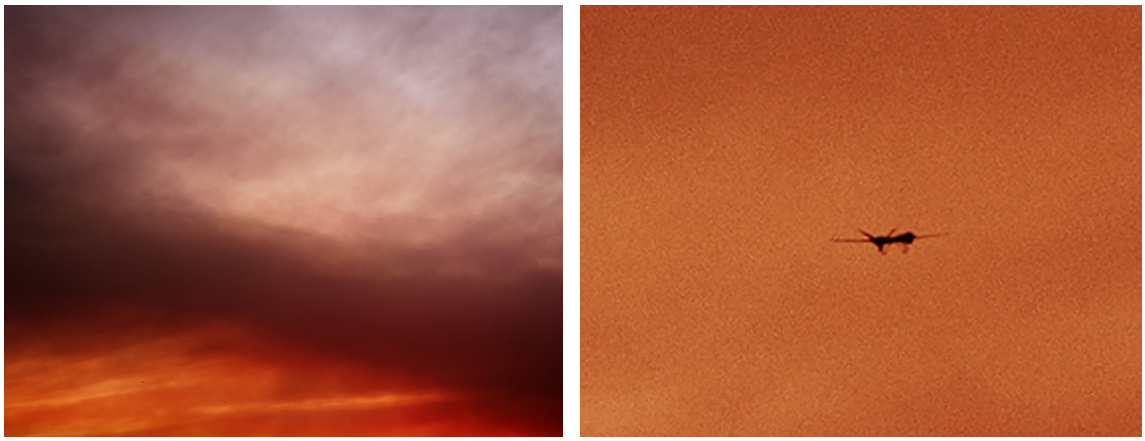
Trevor Paglen, at left, entering an fMRI machine as part of an experiment at Cornell University; at right, an edited compilation of images shown to Dr. Paglen in the machine.
Credit: Paglen Studio (L) and Kuceyeski Lab (R)

In art school, I was taught to read images deliberately and slowly in the same way we might read and critique bits of literature or morsels of philosophy. In the fMRI machine, however, an image is just a visual stimulus to elicit a neurological response.

On my second day of scanning, I learned that my amygdala is highly responsive to a certain type of "boring" image: pictures of small airplanes against a big sky.

This rattled me. Amy Kuceyeski, the radiology professor running this experiment, did not know that in my work as an artist, I spend enormous amounts of time tracking and photographing drones, satellites and aircraft. I felt like my mind had been read, that a machine had reached into

my subconscious to unearth something deep, private and intimate about my professional life and personal interests. A boundary separating the recesses of my mind from the world outside had been permeated.



Trevor Paglen, "Untitled (Reaper Drones)," 2009; and a detail showing the drone in the bottom left of the photograph. Credit: Trevor Paglen/Jessica Silverman and Pace Gallery

In some disorders like depression or anxiety, communication between different regions of the brain is hampered. Dr. Kuceyeski is trying to figure out what kind of visual stimuli could reshape brain networks in those regions and help alleviate symptoms. If scientists knew what someone's brain looked like while they were experiencing anxiety, could mental health clinicians generate a specific image that could help bring them out of it? Other brain researchers are pursuing similar work to help people with language disorders, monitor neurological conditions and diagnose mental health disorders with greater precision.

But the same technology that is accelerating this work — artificial intelligence — is speeding us into a different moment as well.

Modern-day neuroscience isn't yet trying to control anyone's mind. But several labs are pursuing some form of mind-reading, even if they wouldn't call it that. In a 2011 experiment, Jack Gallant at the University of California, Berkeley, showed that if you record people's brain activity while they watch a movie, you can use clever programming combined with computer modeling (basically, a crude precursor to an A.I. system) to reconstruct a blurry, ghostly approximation of what they saw. Since that breakthrough, A.I. has gotten better at recovering higher-fidelity images from brain activity — each one a little sharper, and a little more unnerving, than the last.



A video demo of Jack Gallant's 2011 experiment. On the left are the original videos shown to participants; on the right are reproductions produced by computer modeling based on participants' measured brain activity.
Credit: Shinji Nishimoto and the Gallant lab at UC Berkeley

Alex Huth, a former student of Dr. Gallant's who is studying how the brain processes language, has spent the past several years building systems that reconstruct language from brain activity. In one demo, a participant silently narrates, using only thoughts, a story being heard on a podcast while a screen shows an A.I. model textually reconstructing the narrative, based on the participant's fMRI data. In another, Dr. Huth watches clips of an animated video of a dragon in the scanner, and A.I. uses his neural activity to describe through text the action in the clip on a second monitor.

Large language models allow researchers to transform noisy, coarse brain signals into coherent text. If a neural sensor detects a subject thinking something like "quick … jumps … lazy," a large language model can infer the full thought: "The quick brown fox jumps over the lazy dog."

Dr. Huth cautions that this is just a superficial interpretation of signals from the brain. In other words, brain reading, not mind reading. And it's prone to errors. But he acknowledges the technology may one day take us to undesired places. Imagine, he said, a workplace headset that lets you control your computer using thought alone, but also flags any thoughts related to, say, unionizing. "This is bad. Really deeply bad — and also kind of realistically bad."

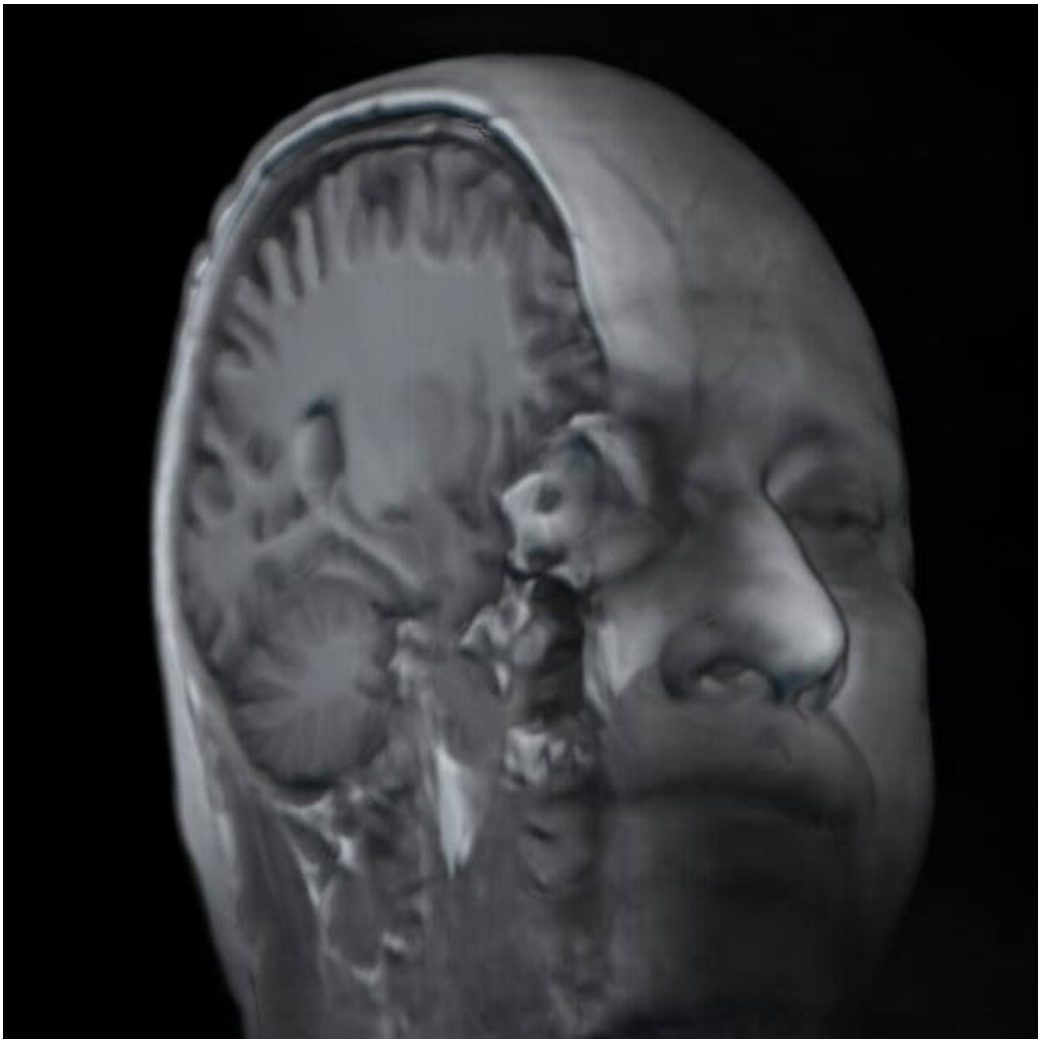


A 2023 experiment run using MinD-Video, an A.I. model, shows how far the technology has progressed. The left panel shows the natural-video stimulus presented to a participant during fMRI scanning; the right panel shows the corresponding video generated by MinD-Video from the participant's fMRI measurements.
Credit: MinD-Video (Zijiao Chen, Jiaxin Qing, and Juan Helen Zhou; NeurIPS 2023)

If you can read the brain, you might eventually be able to write to it as well. At Columbia University, the neurobiologist Rafael Yuste, who is studying how the brain perceives things, has shown that you can point lasers at the brains of mice and make them "see" things that aren't there. "The cortex of the mouse is just like the cortex of a human. Whatever we can do in a mouse today, you can do in a human tomorrow. It's just a matter of time," said Dr. Yuste. He called this breakthrough his Oppenheimer moment, filling him with dread he hadn't felt before. A McGill University neuroscientist, Pouya Bashivan, took these findings a step further, when his group built A.I. models that could accurately predict not just how a monkey's brain would respond to any given image, but also what kind of image could activate the brain in a specific, prescribed way.

That finding is what led me to Dr. Kuceyeski's lab at Cornell, which sought to replicate the same thing in humans. The lab has built a system called NeuroGen, consists of an A.I. model that can reasonably predict how a person's brain will respond when exposed to a specific image. It can

also determine what image would produce a specific pattern of brain activity, and then use generative A.I. to actually make that image.



A grayscale reconstruction of Trevor Paglen's brain from fMRI scans, showing an interior slice of his brain and the blood flow throughout his cranium.
Credit: Kuceyeski Lab, via Paglen Studio

All of this work is cascading at the same time that neurotechnology start-ups, flush with new investments, are proliferating. Dr. Bashivan imagines people one day being able to control their own moods, and dial their emotions up or down as if with a thermostat.

Meta paid between $500 million and $1 billion to acquire a start-up whose wristband translates brain signals into computer commands. Microsoft holds a patent on a possible electroencephalogram (EEG) device that lets users navigate apps with their brains. Apple has signaled it may integrate EEG sensors into AirPods.

We already live in a world of mass surveillance, ruthless competition and corporations determined to monetize every last bit of microdata about us that they can. It's not a stretch to imagine such devices ushering in a world where custom advertisements are generated for your brain; where police interrogation extends to neurological monitors; where the Netflix show you're watching adjusts its plot in real time to suit your neural architecture.

We also live in a world of private joys and daydreams, where our minds wander. Those things are worth protecting, and even researchers in this space know that. Dr. Yuste has co-founded the Neurorights Foundation to build legal guardrails around brain data.