The Science of Mind Reading

Researchers are pursuing age-old questions about the nature of thoughts—and learning how to read them.

By James Somers November 29, 2021

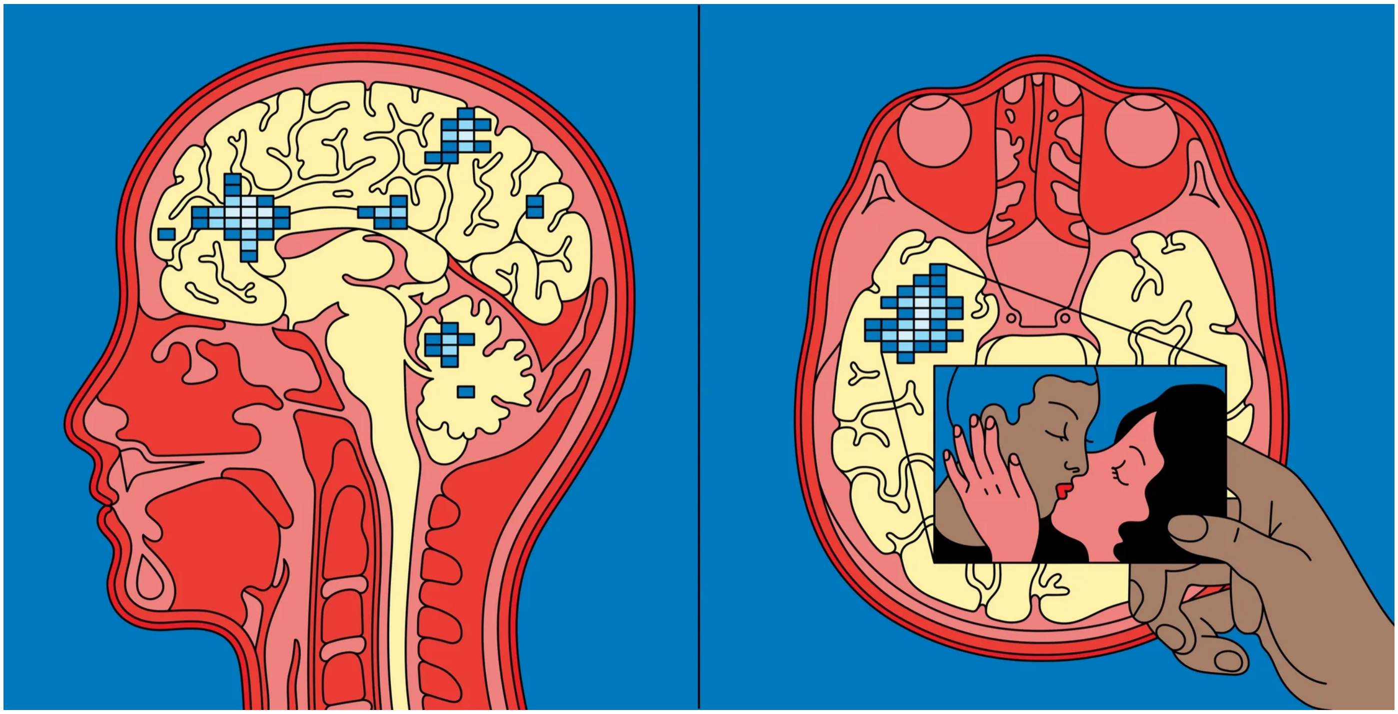

It isn't so much that brain scans have improved—it's that we've got better at reading them. ***Illustration by Laura Edelbacher***

Content

This content can also be viewed on the site it originates from.

One night in October, 2009, a young man lay in an fMRI scanner in Liège, Belgium. Five years earlier, he'd suffered a head trauma in a motorcycle accident, and since then he hadn't spoken. He was said to be in a "vegetative state." A neuroscientist named Martin Monti sat in the next room, along with a few other researchers. For years, Monti and his postdoctoral adviser, Adrian Owen, had been studying vegetative patients, and they had developed two controversial hypotheses. First, they believed that someone could lose the ability to move or even blink while still being conscious; second, they thought that they had devised a method for communicating with such “locked-in” people by detecting their unspoken thoughts.

In a sense, their strategy was simple. Neurons use oxygen, which is carried through the bloodstream inside molecules of hemoglobin. Hemoglobin contains iron, and, by tracking the iron, the magnets in fMRI machines can build maps of brain activity. Picking out signs of consciousness amid the swirl seemed nearly impossible. But, through trial and error, Owen’s group had devised a clever protocol. They’d discovered that if a person imagined walking around her house there was a spike of activity in her parahippocampal gyrus—a finger-shaped area buried deep in the temporal lobe. Imagining playing tennis, by contrast, activated the premotor cortex, which sits on a ridge near the skull. The activity was clear enough to be seen in real time with an fMRI machine. In a 2006 study published in the journal *Science*, the researchers reported that they had asked a locked-in person to think about tennis, and seen, on her brain scan, that she had done so.

With the young man, known as Patient 23, Monti and Owen were taking a further step: attempting to have a conversation. They would pose a question and tell him that he could signal “yes” by imagining playing tennis, or “no” by thinking about walking around his house. In the scanner control room, a monitor displayed a cross-section of Patient 23’s brain. As different areas consumed blood oxygen, they shimmered red, then bright orange. Monti knew where to look to spot the yes and the no signals.

He switched on the intercom and explained the system to Patient 23. Then he asked the first question: “Is your father’s name Alexander?”

The man’s premotor cortex lit up. He was thinking about tennis—yes.

“Is your father’s name Thomas?”

Activity in the parahippocampal gyrus. He was imagining walking around his house—no.

“Do you have any brothers?”

Tennis—yes.

“Do you have any sisters?”

House—no.

“Before your injury, was your last vacation in the United States?”

Tennis—yes.

The answers were correct. Astonished, Monti called Owen, who was away at a conference. Owen thought that they should ask more questions. The group ran through some possibilities. “Do you like pizza?” was dismissed as being too imprecise. They decided to probe more deeply. Monti turned the intercom back on.

“Do you want to die?” he asked.

Cartoon by Liana Finck

For the first time that night, there was no clear answer.

That winter, the results of the study were published in *The New England Journal of Medicine*. The paper caused a sensation. The Los Angeles *Times* wrote a story about it, with the headline “*Brains of Vegetative Patients Show Life*.” Owen eventually estimated that twenty per cent of patients who were presumed to be vegetative were actually awake. This was a discovery of enormous practical consequence: in subsequent years, through painstaking fMRI sessions, Owen's group found many patients who could interact with loved ones and answer questions about their own care. The conversations improved their odds of recovery. Still, from a purely scientific perspective, there was something unsatisfying about the method that Monti and Owen had developed with Patient 23. Although they had used the words "tennis" and "house" in communicating with him, they'd had no way of knowing for sure that he was thinking about those specific things. They had been able to say only that, in response to those prompts, thinking was happening in the associated brain areas. "Whether the person was imagining playing tennis, football, hockey, swimming—we don't know," Monti told me recently.

During the past few decades, the state of neuroscientific mind reading has advanced substantially. Cognitive psychologists armed with an fMRI machine can tell whether a person is having depressive thoughts; they can see which concepts a student has mastered by comparing his brain patterns with those of his teacher. By analyzing brain scans, a computer system can edit together crude reconstructions of movie clips you've watched. One research group has used similar technology to accurately describe the dreams of sleeping subjects. In another lab, scientists have scanned the brains of people who are reading the J. D. Salinger short story "Pretty Mouth and Green My Eyes," in which it is unclear until the end whether or not a character is having an affair. From brain scans alone, the researchers can tell which interpretation readers are leaning toward, and watch as they change their minds.

I first heard about these studies from Ken Norman, the fifty-year-old chair of the psychology department at Princeton University and an expert on thought decoding. Norman works at the Princeton Neuroscience Institute, which is housed in a glass structure, constructed in 2013, that spills over a low hill on the south side of campus. P.N.I. was conceived as a center where psychologists, neuroscientists, and computer scientists could blend their approaches to studying the mind; M.I.T. and Stanford have invested in similar cross-disciplinary institutes. At P.N.I., undergraduates still participate in old-school psych experiments involving surveys and flash cards. But upstairs, in a lab that studies child development, toddlers wear tiny hats outfitted with infrared brain scanners, and in the basement the skulls of genetically engineered mice are sliced open, allowing individual neurons to be controlled with lasers. A server room with its own high-performance computing cluster analyzes the data generated from these experiments.

Norman, whose jovial intelligence and unruly beard give him the air of a high-school science teacher, occupies an office on the ground floor, with a view of a grassy field. The bookshelves behind his desk contain the intellectual DNA of the institute, with William James next to texts on machine learning. Norman explained that fMRI machines hadn't advanced that much; instead, artificial intelligence had transformed how scientists read neural data. This had helped shed light on an ancient philosophical mystery. For centuries, scientists had dreamed of locating thought inside the head but had run up against the vexing question of what it means for thoughts to exist in physical space. When Erasistratus, an ancient Greek anatomist, dissected the brain, he suspected that its many folds were the key to intelligence, but he could not say how thoughts were packed into the convoluted mass. In the seventeenth century, Descartes suggested that mental life arose in the pineal gland, but he didn't have a good theory of what might be found there. Our mental worlds contain everything from the taste of bad wine to the idea of bad taste. How can so many thoughts nestle within a few pounds of tissue?

Now, Norman explained, researchers had developed a mathematical way of understanding thoughts. Drawing on insights from machine learning, they conceived of thoughts as collections of points in a dense "meaning space." They could see how these points were interrelated and encoded by neurons. By cracking the code, they were beginning to produce an inventory of the mind. "The space of possible thoughts that people can think is big—but it's not infinitely big," Norman said. A detailed map of the concepts in our minds might soon be within reach.

Norman invited me to watch an experiment in thought decoding. A postdoctoral student named Manoj Kumar led us into a locked basement lab at P.N.I., where a young woman was lying in the tube of an fMRI scanner. A screen mounted a few inches above her face played a slide show of stock images: an empty beach, a cave, a forest.

"We want to get the brain patterns that are associated with different subclasses of scenes," Norman said.

As the woman watched the slide show, the scanner tracked patterns of activation among her neurons. These patterns would be analyzed in terms of "voxels"—areas of activation that are roughly a cubic millimetre in size. In some ways, the fMRI data was extremely coarse: each voxel represented the oxygen consumption of about a million neurons, and could be updated only every few seconds, significantly more slowly than neurons fire. But, Norman said, "it turned out that that information was in the data we were collecting—we just weren't being as smart as we possibly could about how we'd churn through that data." The breakthrough came when researchers figured out how to track patterns playing out across tens of thousands of voxels at a time, as though each were a key on a piano, and thoughts were chords.

The origins of this approach, I learned, dated back nearly seventy years, to the work of a psychologist named Charles Osgood. When he was a kid, Osgood received a copy of Roget's Thesaurus as a gift. Poring over the book, Osgood recalled, he formed a "vivid image of words as clusters of starlike points in an immense space." In his postgraduate days, when his colleagues were debating how cognition could be shaped by culture, Osgood thought back on this image. He wondered if, using the idea of "semantic space," it might be possible to map the differences among various styles of thinking.

Osgood conducted an experiment. He asked people to rate twenty concepts on fifty different scales. The concepts ranged widely: *BOULDER, ME, TORNADO, MOTHER*. So did the scales, which were defined by opposites: fair-unfair, hot-cold, fragrant-foul. Some ratings were difficult: is a *TORNADO* fragrant or foul? But the idea was that the method would reveal fine and even elusive shades of similarity and difference among concepts. "Most English-speaking Americans feel that there is a difference, somehow, between 'good' and 'nice' but find it difficult to explain," Osgood wrote. His surveys found that, at least for nineteen-fifties college students, the two concepts overlapped much of the time. They diverged for nouns that had a male or female slant. MOTHER might be rated nice but not good, and COP vice versa. Osgood concluded that "good" was "somewhat stronger, rougher, more angular, and larger" than "nice."

Osgood became known not for the results of his surveys but for the method he invented to analyze them. He began by arranging his data in an imaginary space with fifty dimensions—one for fair-unfair, a second for hot-cold, a third for fragrant-foul, and so on. Any given concept, like *TORNADO*, had a rating on each dimension—and, therefore, was situated in what was known as high-dimensional space. Many concepts had similar locations on multiple axes: kind-cruel and honest-dishonest, for instance. Osgood combined these dimensions. Then he looked for new similarities, and combined dimensions again, in a process called "factor analysis."

When you reduce a sauce, you meld and deepen the essential flavors. Osgood did something similar with factor analysis. Eventually, he was able to map all the concepts onto a space with just three dimensions. The first dimension was “evaluative”—a blend of scales like good-bad, beautiful-ugly, and kind-cruel. The second had to do with “potency”: it consolidated scales like large-small and strong-weak. The third measured how “active” or “passive” a concept was. Osgood could use these three key factors to locate any concept in an abstract space. Ideas with similar coördinates, he argued, were neighbors in meaning.

For decades, Osgood’s technique found modest use in a kind of personality test. Its true potential didn’t emerge until the nineteen-eighties, when researchers at Bell Labs were trying to solve what they called the “vocabulary problem.” People tend to employ lots of names for the same thing. This was an obstacle for computer users, who accessed programs by typing words on a command line. George Furnas, who worked in the organization’s human-computer-interaction group, described using the company’s internal phone book. “You’re in your office, at Bell Labs, and someone has stolen your calculator,” he said. “You start putting in ‘police,’ or ‘support,’ or ‘theft,’ and it doesn’t give you what you want. Finally, you put in ‘security,’ and it gives you that. But it actually gives you two things: something about the Bell Savings and Security Plan, and also the thing you’re looking for.” Furnas’s group wanted to automate the finding of synonyms for commands and search terms.

They updated Osgood’s approach. Instead of surveying undergraduates, they used computers to analyze the words in about two thousand technical reports. The reports themselves—on topics ranging from graph theory to user-interface design—suggested the dimensions of the space; when multiple reports used similar groups of words, their dimensions could be combined. In the end, the Bell Labs researchers made a space that was more complex than Osgood’s. It had a few hundred dimensions. Many of these dimensions described abstract or “latent” qualities that the words had in common—connections that wouldn’t be apparent to most English speakers. The researchers called their technique “latent semantic analysis,” or L.S.A.

At first, Bell Labs used L.S.A. to create a better internal search engine. Then, in 1997, Susan Dumais, one of Furnas’s colleagues, collaborated with a Bell Labs cognitive scientist, Thomas Landauer, to develop an A.I. system based on it. After processing Grolier’s American Academic Encyclopedia, a work intended for young students, the A.I. scored respectably on the multiple-choice Test of English as a Foreign Language. That year, the two researchers co-wrote a paper that addressed the question “How do people know as much as they do with as little information as they get?” They suggested that our minds might use something like L.S.A., making sense of the world by reducing it to its most important differences and similarities, and employing this distilled knowledge to understand new things. Watching a Disney movie, for instance, I immediately identify a character as “the bad guy”: Scar, from “The Lion King,” and Jafar, from “Aladdin,” just seem close together. Perhaps my brain uses factor analysis to distill thousands of attributes—height, fashion sense, tone of voice—into a single point in an abstract space. The perception of bad-guy-ness becomes a matter of proximity.

In the following years, scientists applied L.S.A. to ever-larger data sets. In 2013, researchers at Google unleashed a descendant of it onto the text of the whole World Wide Web. Google’s algorithm turned each word into a “vector,” or point, in high-dimensional space. The vectors generated by the researchers’ program, word2vec, are eerily accurate: if you take the vector for “king” and subtract the vector for “man,” then add the vector for “woman,” the closest nearby vector is “queen.” Word vectors became the basis of a much improved Google Translate, and enabled the auto-completion of sentences in Gmail. Other companies, including Apple and

Amazon, built similar systems. Eventually, researchers realized that the "vectorization" made popular by L.S.A. and word2vec could be used to map all sorts of things. Today's facial-recognition systems have dimensions that represent the length of the nose and the curl of the lips, and faces are described using a string of coördinates in "face space." Chess A.I.s use a similar trick to "vectorize" positions on the board. The technique has become so central to the field of artificial intelligence that, in 2017, a new, hundred-and-thirty-five-million-dollar A.I. research center in Toronto was named the Vector Institute. Matthew Botvinick, a professor at Princeton whose lab was across the hall from Norman's, and who is now the head of neuroscience at DeepMind, Alphabet's A.I. subsidiary, told me that distilling relevant similarities and differences into vectors was "the secret sauce underlying all of these A.I. advances."

In 2001, a scientist named Jim Haxby brought machine learning to brain imaging: he realized that voxels of neural activity could serve as dimensions in a kind of thought space. Haxby went on to work at Princeton, where he collaborated with Norman. The two scientists, together with other researchers, concluded that just a few hundred dimensions were sufficient to capture the shades of similarity and difference in most fMRI data. At the Princeton lab, the young woman watched the slide show in the scanner. With each new image—beach, cave, forest—her neurons fired in a new pattern. These patterns would be recorded as voxels, then processed by software and transformed into vectors. The images had been chosen because their vectors would end up far apart from one another: they were good landmarks for making a map. Watching the images, my mind was taking a trip through thought space, too.

The larger goal of thought decoding is to understand how our brains mirror the world. To this end, researchers have sought to watch as the same experiences affect many people's minds simultaneously. Norman told me that his Princeton colleague Uri Hasson has found movies especially useful in this regard. They “pull people’s brains through thought space in synch,” Norman said. “What makes Alfred Hitchcock the master of suspense is that all the people who are watching the movie are having their brains yanked in unison. It’s like mind control in the literal sense.”

One afternoon, I sat in on Norman’s undergraduate class “fMRI Decoding: Reading Minds Using Brain Scans.” As students filed into the auditorium, setting their laptops and water bottles on tables, Norman entered wearing tortoiseshell glasses and earphones, his hair dishevelled.

He had the class watch a clip from “Seinfeld” in which George, Susan (an N.B.C. executive he is courting), and Kramer are hanging out with Jerry in his apartment. The phone rings, and Jerry answers: it’s a telemarketer. Jerry hangs up, to cheers from the studio audience.

“Where was the event boundary in the clip?” Norman asked. The students yelled out in chorus, “When the phone rang!” Psychologists have long known that our minds divide experiences into segments; in this case, it was the phone call that caused the division.

Norman showed the class a series of slides. One described a 2017 study by Christopher Baldassano, one of his postdocs, in which people watched an episode of the BBC show “Sherlock” while in an fMRI scanner. Baldassano’s guess going into the study was that some voxel patterns would be in constant flux as the video streamed—for instance, the ones involved in color processing. Others would be more stable, such as those representing a character in the show. The study confirmed these predictions. But Baldassano also found groups of voxels that held a stable pattern throughout each scene, then switched when it was over. He concluded that these constituted the scenes’ voxel “signatures.”

Norman described another study, by Asieh Zadbood, in which subjects were asked to narrate “Sherlock” scenes—which they had watched earlier—aloud. The audio was played to a second group, who’d never seen the show. It turned out that no matter whether someone watched a scene, described it, or heard about it, the same voxel patterns recurred. The scenes existed independently of the show, as concepts in people’s minds.

Through decades of experimental work, Norman told me later, psychologists have established the importance of scripts and scenes to our intelligence. Walking into a room, you might forget why you came in; this happens, researchers say, because passing through the doorway brings one mental scene to a close and opens another. Conversely, while navigating a new airport, a “getting to the plane” script knits different scenes together: first the ticket counter, then the security line, then the gate, then the aisle, then your seat. And yet, until recently, it wasn’t clear what you’d find if you went looking for “scripts” and “scenes” in the brain.

In a recent P.N.I. study, Norman said, people in an fMRI scanner watched various movie clips of characters in airports. No matter the particulars of each clip, the subjects’ brains all shimmered through the same series of events, in keeping with boundary-defining moments that any of us would recognize. The scripts and the scenes were real—it was possible to detect them with a machine. What most interests Norman now is how they are learned in the first place. How do we identify the scenes in a story? When we enter a strange airport, how do we know intuitively where to look for the security line? The extraordinary difficulty of such feats is obscured by how easy they feel—it’s rare to be confused about how to make sense of the world. But at some point everything was new. When I was a toddler, my parents must have taken me to the supermarket for the first time; the fact that, today, all supermarkets are somehow familiar dims the strangeness of that experience. When I was learning to drive, it was overwhelming: each intersection and lane change seemed chaotic in its own way. Now I hardly have to think about them. My mind instantly factors out all but the important differences.

Norman clicked through the last of his slides. Afterward, a few students wandered over to the lectern, hoping for an audience with him. For the rest of us, the scene was over. We packed up, climbed the stairs, and walked into the afternoon sun.

Like Monti and Owen with Patient 23, today's thought-decoding researchers mostly look for specific thoughts that have been defined in advance. But a "general-purpose thought decoder," Norman told me, is the next logical step for the research. Such a device could speak aloud a person's thoughts, even if those thoughts have never been observed in an fMRI machine. In 2018, Botvinick, Norman's hall mate, helped write a paper in the journal *Nature Communications* titled "Toward a Universal Decoder of Linguistic Meaning from Brain Activation." A team of researchers led by Botvinick's former postdoc, Francisco Pereira, and Evelina Fedorenko, a neuroscientist at M.I.T., had built a primitive form of what Norman described: a system that could decode novel sentences that subjects read silently to themselves. The system learned which brain patterns were evoked by certain words, and used that knowledge to guess which words were implied by the new patterns it encountered.

"I'd love a pet right now, but I travel too much."

Cartoon by Liza Donnelly

The work at Princeton was funded by iARPA, an R. & D. organization that's run by the Office of the Director of National Intelligence. Brandon Minnery, the iARPA project manager for the Knowledge Representation in Neural Systems program at the time, told me that he had some applications in mind. If you knew how knowledge was represented in the brain, you might be able to distinguish between novice and expert intelligence agents. You might learn how to teach languages more effectively by seeing how closely a student's mental representation of a word matches that of a native speaker. Minnery's most fanciful idea—"Never an official focus of the program," he said—was to change how databases are indexed. Instead of labelling items by hand, you could show an item to someone sitting in an fMRI scanner—the person's brain state could be the label. Later, to query the database, someone else could sit in the scanner and simply think of whatever she wanted. The software could compare the searcher's brain state with the indexer's. It would be the ultimate solution to the vocabulary problem.

Jack Gallant, a professor at Berkeley who has used thought decoding to reconstruct video montages from brain scans—as you watch a video in the scanner, the system pulls up frames from similar YouTube clips, based only on your voxel patterns—suggested that one group of people interested in decoding were Silicon Valley investors. "A future technology would be a portable hat—like a thinking hat," he said. He imagined a company paying people thirty thousand dollars a year to wear the thinking hat, along with video-recording eyeglasses and other sensors, allowing the system to record everything they see, hear, and think, ultimately creating an exhaustive inventory of the mind. Wearing the thinking hat, you could ask your computer a question just by imagining the words. Instantaneous translation might be possible. In theory, a pair of wearers could skip language altogether, conversing directly, mind to mind. Perhaps we could even communicate across species. Among the challenges the designers of such a system would face, of course, is the fact that today's fMRI machines can weigh more than twenty thousand pounds. There are efforts under way to make powerful miniature imaging devices, using lasers, ultrasound, or even microwaves. "It's going to require some sort of punctuated-equilibrium technology revolution," Gallant said. Still, the conceptual foundation, which goes back to the nineteen-fifties, has been laid.

Recently, I asked Owen what the new thought-decoding technology meant for locked-in patients. Were they close to having fluent conversations using something like the general-purpose thought decoder? "Most of that stuff is group studies in healthy participants," Owen told me. "The really tricky problem is doing it in a single person. Can you get robust enough data?" Their bare-bones protocol—thinking about tennis equals yes; thinking about walking around the house equals no—relied on straightforward signals that were statistically robust. It turns out that the same protocol, combined with a series of yes-or-no questions ("Is the pain in the lower half of your body? On the left side?"), still works best. "Even if you could do it, it would take longer to decode them saying 'it is in my right foot' than to go through a simple series of yes-or-no questions," Owen said. "For the most part, I'm quietly sitting and waiting. I have no doubt that, some point down the line, we will be able to read minds. People will be able to articulate, 'My name is Adrian, and I'm British,' and we'll be able to decode that from their brain. I don't think it's going to happen in probably less than twenty years."

In some ways, the story of thought decoding is reminiscent of the history of our understanding of the gene. For about a hundred years after the publication of Charles Darwin's "On the Origin of Species," in 1859, the gene was an abstraction, understood only as something through which traits passed from parent to child. As late as the nineteen-fifties, biologists were still asking what, exactly, a gene was made of. When James Watson and Francis Crick finally found the double helix, in 1953, it became clear how genes took physical form. Fifty years later, we could sequence the human genome; today, we can edit it.

Thoughts have been an abstraction for far longer. But now we know what they really are: patterns of neural activation that correspond to points in meaning space. The mind—the only truly private place—has become inspectable from the outside. In the future, a therapist, wanting to understand how your relationships run awry, might examine the dimensions of the patterns your brain falls into. Some epileptic patients about to undergo surgery have intracranial probes put into their brains; researchers can now use these probes to help steer the patients' neural patterns away from those associated with depression. With more fine-grained control, a mind could be driven wherever one liked. (The imagination reels at the possibilities, for both good and ill.) Of course, we already do this by thinking, reading, watching, talking—actions that, after I'd learned about thought decoding, struck me as oddly concrete. I could picture the patterns of my thoughts flickering inside my mind. Versions of them are now flickering in yours.

On one of my last visits to Princeton, Norman and I had lunch at a Japanese restaurant called Ajiten. We sat at a counter and went through the familiar script. The menus arrived; we looked them over. Norman noticed a dish he hadn't seen before—"a new point in ramen space," he said. Any minute now, a waiter was going to interrupt politely to ask if we were ready to order.

"You have to carve the world at its joints, and figure out: what are the situations that exist, and how do these situations work?" Norman said, while jazz played in the background. "And that's a very complicated problem. It's not like you're instructed that the world has fifteen different ways of being, and here they are!" He laughed. "When you're out in the world, you have to try to infer what situation you're in." We were in the lunch-at-a-Japanese-restaurant situation. I had never been to this particular restaurant, but nothing about it surprised me. This, it turns out, might be one of the highest accomplishments in nature.

Norman told me that a former student of his, Sam Gershman, likes using the terms "lumping" and "splitting" to describe how the mind's meaning space evolves. When you encounter a new stimulus, do you lump it with a concept that's familiar, or do you split off a new concept? When navigating a new airport, we lump its metal detector with those we've seen before, even if this one is a different model, color, and size. By contrast, the first time we raised our hands inside a millimetre-wave scanner—the device that has replaced the walk-through metal detector—we split off a new category.

Norman turned to how thought decoding fit into the larger story of the study of the mind. "I think we're at a point in cognitive neuroscience where we understand a lot of the pieces of the puzzle," he said. The cerebral cortex—a crumply sheet laid atop the rest of the brain—warps and compresses experience, emphasizing what's important. It's in constant communication with other brain areas, including the hippocampus, a seahorse-shaped structure in the inner part of the temporal lobe. For years, the hippocampus was known only as the seat of memory; patients who'd had theirs removed lived in a perpetual present. Now we were seeing that the hippocampus stores summaries provided to it by the cortex: the sauce after it's been reduced. We cope with reality by building a vast library of experience—but experience that has been distilled along the dimensions that matter. Norman's research group has used fMRI technology to find voxel patterns in the cortex that are reflected in the hippocampus. Perhaps the brain is like a hiker comparing the map with the territory.

In the past few years, Norman told me, artificial neural networks that included basic models of both brain regions had proved surprisingly powerful. There was a feedback loop between the study of A.I. and the study of the real human mind, and it was getting faster. Theories about human memory were informing new designs for A.I. systems, and those systems, in turn, were suggesting ideas about what to look for in real human brains. "It's kind of amazing to have gotten to this point," he said.

On the walk back to campus, Norman pointed out the Princeton University Art Museum. It was a treasure, he told me.

“What’s in there?” I asked.

“Great art!” he said

After we parted ways, I returned to the museum. I went to the downstairs gallery, which contains artifacts from the ancient world. Nothing in particular grabbed me until I saw a West African hunter’s tunic. It was made of cotton dyed the color of dark leather. There were teeth hanging from it, and claws, and a turtle shell—talismans from past kills. It struck me, and I lingered for a moment before moving on.

Six months later, I went with some friends to a small house in upstate New York. On the wall, out of the corner of my eye, I noticed what looked like a blanket—a kind of fringed, hanging decoration made of wool and feathers. It had an odd shape; it seemed to pull toward something I’d seen before. I stared at it blankly. Then came a moment of recognition, along dimensions I couldn’t articulate—more active than passive, partway between alive and dead. There, the chest. There, the shoulders. The blanket and the tunic were distinct in every way, but somehow still neighbors. My mind had split, then lumped. Some voxels had shimmered. In the vast meaning space inside my head, a tiny piece of the world was finding its proper place. ♦

This article has been updated to include two of the lead researchers who built a system that could decode novel sentences that subjects read silently to themselves.

New Yorker Favorites

- How we became infected by chain e-mail.
- Twelve classic movies to watch with your kids.
- The secret lives of fungi.

- The photographer who claimed to capture the ghost of Abraham Lincoln.
- Why are Americans still uncomfortable with atheism?
- The enduring romance of the night train.
- Sign up for our daily newsletter to receive the best stories from *The New Yorker*.