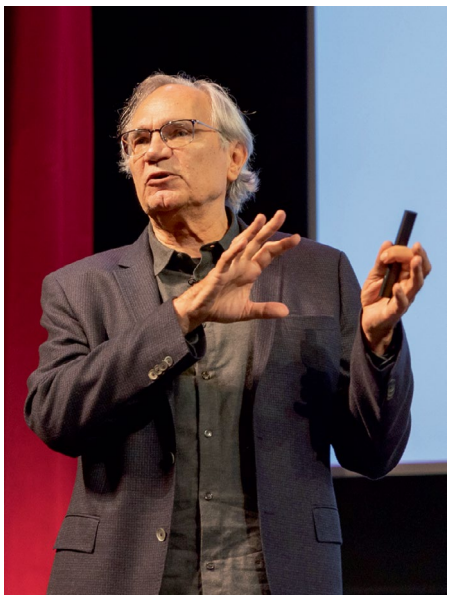
https://doi.org/10.1038/s41593-026-02450-6

In conversation with Richard Andersen

Check for updates

As part of our special issue on neurotechnology, we spoke with *Richard Andersen*, James G. Boswell Professor of Neuroscience at Caltech, about the discoveries that helped to establish the posterior parietal cortex as a key hub for intention and action planning. He looks back on decades of discoveries in the posterior parietal cortex and discusses how advances in neural recording, decoding and computation are reshaping our understanding of cognition.



You've had a fascinating career spanning a particularly transformative period in neuroscience, from single-neuron recordings through to AI-based decoding. Looking back, what technological advances do you think have most changed the kinds of questions we can ask about the brain?

Early on, when I was looking for a postdoc, awake recordings in non-human primates were still relatively rare. At that point, the main places doing that work were NIH and Johns Hopkins. At NIH, Edward Evarts was studying the motor system, and Robert Wurtz and his group were working on vision. At Johns Hopkins, Vernon Mountcastle was recording from the posterior parietal cortex, which I was very interested in as it seemed both fascinating and mysterious. Mountcastle had found neurons that appeared to encode eye movements, reaching, and grasping[1], so the area looked much more motor-related than people had previously appreciated.

About two years later, Mickey Goldberg and colleagues at NIH, coming from Wurtz's group, published papers that challenged Mountcastle's interpretation[2,3]. Their argument was that these effects might be visual or attentional, rather than motor. So when I went to do my postdoc with Mountcastle, one of the big questions was whether these signals were motor plans or simply related to visual processing and attention.

Both views contained part of the truth. There were certainly sensory influences, but there were also motor-related signals. What we found was that the activity wasn't really coding the final movement itself, as Mountcastle had thought. Instead, it seemed to encode the animal's intention or plan.

What really clinched that idea was a paper Larry Snyder, Aaron Batista and I published in *Nature* in 1997[4]. Because monkeys can separate eye movements from reaching movements, we could ask them either to plan a saccade or to plan a reach. We found that the lateral intraparietal area, for eye movements, became active specifically when the animal was planning a saccade, while the reach-related region, now often called the medial intraparietal area, became active specifically when the animal was planning a reach. If the activity had simply reflected attention or sensory processing, it shouldn't have depended on what action the animal intended to make. That was one major leap forward.

Another big advance was the development of techniques for recording from populations of neurons. If you look only at single-unit activity, it's often difficult to understand exactly what's happening. But when you record from populations of cells, you start to see that each neuron contributes a small piece of the puzzle. Taken together, the underlying computation becomes much clearer. These advances also enabled us, and many other groups, to make the transition to recording in humans. That's really where today's brain–computer interfaces come from. It has opened up entirely new possibilities, both for clinical applications and for doing fundamental neuroscience directly in humans.

So I would say those were two of the biggest technological advances: first, the development of sophisticated awake behaving animal recordings that allowed us to identify intention-related signals, and second, large-scale population recording, which ultimately led to modern brain–computer interfaces and human neural recording.

You mentioned the 1997 *Nature* paper as a key result. How much did that study change the field's thinking about the posterior parietal cortex, and what evidence over the years has convinced you that the framework has held up?

We're still one of the few groups that records extensively from the posterior parietal cortex. Most other groups have concentrated on motor cortex, which we also record from. Being able to compare motor cortex and posterior parietal cortex directly has been very useful for understanding the different roles they have. In monkeys, movement plans are anatomically separated within the intraparietal sulcus. In addition to the regions described above, Sakata and colleagues[5] described another important region that was grasp-related, the anterior intraparietal area (AIP). Those findings have stood the test of time. Right now, we're using functional ultrasound, which we find can be extremely powerful. Once again, we see this separation of function. We can then use Neuropixels probes to record within these areas in much greater detail, and the same basic organization continues to hold up.

In humans, the story is more complicated. Before implantation, we use fMRI while participants imagine or attempt grasping, even if paralyzed, to identify the human homologue of AIP. But when we implant in this region, we find an enormous amount of mixed activity. It's reminiscent of the mixed selectivity described by Earl Miller and Stefano Fusi[6], but taken to an extreme. With a single 4 × 4-mm, 96-electrode array in posterior parietal cortex, we can decode intended movements involving virtually all parts of the body. The signals

CREDIT: CALTECH

are highly mixed even for single neurons. By contrast, signals from implants in motor cortex are much less mixed and look more like what you would expect. Neurons in hand areas tend to code primarily for reaching or hand movements, involving either the ipsilateral or contralateral hand.

The difference could reflect species differences, the specific areas being sampled, or the way the experiments are performed. In monkeys, animals may train on a task for months, whereas humans can perform many new tasks immediately, allowing us to probe a much broader range of behaviors.

How do you view the Intentional Maps framework today? Which aspects do you think have stood the test of time, and are there any parts that need to be revised?

I think it has held up remarkably well in monkeys, and it's generated a great deal of subsequent work. For example, Michael Shadlen and Bill Newsome looked at decision-making in lateral intraparietal regions and found activity very much in line with what we described[7]. What they added was a model for how competing potential action plans can be evaluated and ultimately lead to a decision, in their case involving eye movements. A number of groups have replicated and extended it over the years. We have as well.

Humans may turn out to be more complicated, which perhaps isn't surprising. But the framework gave us a way to ask new questions. Knowing which areas are involved in planning from work in non-human primates means we can design tasks specifically aimed at understanding planning-related activity in the human posterior parietal cortex.

Another question concerns the computations that underlie the transformation from spontaneous thought to action. What do you think we still don't understand about that process?

One thing we found was what we called gain fields[8], which seem to have an important role in transforming information between coordinate systems, for example from eye-centered visual coordinates into the coordinate frames needed to guide limb movements. It was difficult to understand this at the level of individual neurons. But when we looked at populations of cells, the picture became much clearer. I was fortunate to be at the Salk Institute at a time when Geoffrey Hinton, David Rumelhart and others had just developed back-propagation, so we were able to model some of these posterior parietal transformations using neural network approaches[9].

Of course, there are many other interesting computational questions. You brought up one of them, which is how you go from a spontaneous thought to an intended plan. One thing that puzzled me was that we could decode intended body movements without knowing where the eyes were looking. We eventually found that some of these coordinate-frame representations can switch depending on the task[10]. In other words, the reference frame itself can be task dependent.

More generally, the posterior parietal cortex appears to construct an internal representation of the body in relation to the external world. It also receives efference copies, which are probably important for state estimation. So in many ways it fits with ideas from control theory about how the brain keeps track of the body's state and uses that information to guide action.

> "the posterior parietal cortex appears to construct an internal representation of the body in relation to the external world."

More recently, we've become interested in mirror neuron-like selectivity. We see both sensory mirroring, where you represent what another person is feeling, and the more classical responses associated with observing another individual's movements. But we also began to notice something else[11]. The posterior parietal cortex seems to be representing the world outside the body and mapping what is happening in the environment. That led us to collaborate with Nancy Kanwisher and Josh Tenenbaum. They have proposed ideas about physical inference engines[12], and we are finding evidence that the same areas we record from also represent interactions between objects in the external world. The emerging picture is that these regions map both the body and the surrounding environment. On that basis, they help to generate decisions, thoughts and intentions that can ultimately be translated into actions.

Your work has highlighted stable intention signals in posterior parietal cortex, while other groups have emphasized neural dynamics and state-space trajectories. How do you see those ideas fitting together?

A lot of our research could be said to be on representations, that is, internal models, tuning curves, somatotopy and so on. However, dynamical systems analysis has also been valuable in seeing how populations of neurons evolve in real time.

For intention, we designed our experiments specifically to study stable intent. We use a delayed-response paradigm in which a target is presented and then removed. There's nothing left for the human or animal to see, but they continue planning the movement until they receive a signal to execute it. Natural movements, of course, occur much more quickly. More recently, we've begun looking at the role of dynamics in the posterior parietal cortex. We have used dynamical modelling to compare activity in posterior parietal and motor cortex.

In motor cortex, the initiation of movement looks very similar to what Krishna Shenoy, Mark Churchland and their colleagues found in monkeys[13]. There is a brief period of rotational dynamics at movement onset. Their proposal was that once the initial conditions are set, the system evolves autonomously. What's interesting is that we can now examine this in humans. In monkeys, movements are often very ballistic, so it's reasonable that the neural activity resembles signals closely related to muscle control.

In our human studies, however, participants can also perform slower, continuous tracking tasks. In those situations, we still see an initial period of rotational dynamics at movement onset, but it is followed by a more stable pattern of activity. Using switching dynamical models, we've shown that at least two distinct dynamical states seem to be involved[14]. One is the transient rotational dynamic associated with movement initiation. The other is a more stable pattern that persists during ongoing behavior. Motor cortex receives external inputs, probably from posterior parietal cortex, as well as sensory feedback. Whereas posterior parietal activity appears relatively stable and representational, motor cortex dynamics are more transient and continually shaped by incoming information.

Do you feel neuroscience is moving beyond increasingly detailed descriptions of neural activity and toward mechanistic explanations of cognition?

Two themes emerging across neuroscience are mixed selectivity, which I mentioned earlier, and compositional coding. In the posterior parietal cortex, many neurons represent combinations of different effectors and actions.

For example, signals related to eye movement plans can be combined with reach plans within the activity of individual neurons. But looking at the activity of populations of neurons, it's as though the brain is using a set of building blocks that can be combined to generate more complex behaviors. In that sense, it's not entirely unlike language, where a limited set of elements can be recombined to produce a vast number of expressions.

We're seeing something similar in the observation experiments. Neural activity seems to be built from components representing things such as self versus other, different types of action, or different body parts. These elements can then be combined in different ways depending on the situation. These findings suggest that cognition and action are built from fundamental neural components, an idea now being explored by several groups.

Your work has shown that intention and internal speech can be decoded from brain activity. As these technologies become more powerful, what does meaningful cognitive privacy look like, and who should get to decide where the limits are?

Ethical considerations are going to become increasingly important. At the moment, most of this research is still focused on medical applications rather than human enhancement or broader questions of privacy. A couple of years ago, we published a study in *Nature Human Behaviour* describing the decoding of internal speech[15]. To be precise, it wasn't the decoding of thought in general, but rather the internal speech people use when talking to themselves. In that case, it wasn't too much of an ethical problem because participants were performing a structured task in which we explicitly asked them to internally rehearse particular words. The Stanford group has reported similar findings from premotor cortex where participants engage in more spontaneous forms of internal speech during a task[16].

Potentially, that does raise privacy concerns. However, one reason I'm relatively optimistic is that these signals appear to depend strongly on task engagement. In our case, if participants simply don't perform the task, the signal largely disappears. I believe the Stanford group also found that internal and overt speech occupy largely the same neural subspace, with subtle differences between them. That suggests people may be able to exert some control over what can be decoded. So while neural decoding raises legitimate questions about mental privacy, there is also evidence that individuals may be able to exert considerable control over their internal dialog and therefore what can be decoded.

If you could define a grand challenge for systems neuroscience over the next decade, what would it be? What is the most important unanswered question about how thoughts become actions?

Well, as you say, that really would be a grand challenge. We're already able to decode speech, which is a relatively high-level aspect of cognition. At the same time, we're finding increasing evidence for semantic representations in the posterior parietal cortex, which perhaps is not too surprising. To understand how language interacts with semantics, how we acquire knowledge of the world and how that knowledge shapes language, I think it's going to be a major area of research in the years ahead.

"To understand how language interacts with semantics, how we acquire knowledge of the world and how that knowledge shapes language, I think it's going to be a major area of research in the years ahead."

You've mentored many scientists who have gone on to become leaders in the field. What qualities distinguish a great scientist from a good one?

You're putting me on the spot. I'll have to think back to my reference letters! I've been fortunate to have many outstanding students and trainees who have gone on to become leaders in the field. In our area, a strong quantitative background is certainly important. That's just the nature of the work. Beyond that, I think ambition matters. Having the passion to really apply yourself. I've also been fortunate in another respect, which is the environment. I'm lucky, and I often tell my students this too: I'm surrounded by exceptionally talented people. Caltech, MIT, Stanford, Harvard and others attract outstanding students. A lot of the credit belongs not just to the mentor, but to the students and the lab culture they build together.

Looking ten years ahead, what outcome would convince you that the brain–computer interface field has genuinely succeeded?

Right now, both our lab and many others can do remarkable things in a research setting. The challenge is translating those demonstrations into products that are genuinely useful in everyday life. I think there are no major intellectual barriers, only practical ones. We need implants that are less invasive, so people are willing to get them, and easier to deploy. A good model would be the cochlear implant. If someone has paralysis, they should be able to receive a neural implant as routinely as patients today receive a cochlear implant for hearing loss.

I think we'll also see growing applications in neuropsychiatric disorders. That's one area where technologies such as functional ultrasound could prove particularly valuable, because they may be able to track slower-changing neural states and provide forms of stimulation that you could use clinically. The drawback of functional ultrasound, like fMRI, is speed. A colleague I work closely with, Mikhail Shapiro at Caltech, has taken time off to co-found a company called Merge, working on sonogenetics[17], which is analogous to optogenetics. The idea is to make cells sensitive to ultrasound through genetic modification, potentially allowing much faster stimulation and recording. Of course, the major challenge is that it would require delivering viral vectors into the human cortex. Still, I think it's a very interesting new technological direction for brain–computer interfaces.

I hope that within the next decade these technologies become practical enough to be used routinely by patients with neurological disorders. That, to me, would be the clearest sign that the field has succeeded.

Interviewed by Henrietta Howells

Published online: 1 October 2026

References

1. Mountcastle, V. B., Lynch, J. C., Georgopoulos, A., Sakata, H. & Acuña, C. *J. Neurophysiol.* **38**, 871–908 (1975).
2. Robinson, D. L., Goldberg, M. E. & Stanton, G. B. *J. Neurophysiol.* **41**, 910–932 (1978).
3. Bushnell, M. C., Goldberg, M. E. & Robinson, D. L. *J. Neurophysiol.* **46**, 755–772 (1981).
4. Snyder, L. H., Batista, A. P. & Andersen, R. A. *Nature* **386**, 167–170 (1997).
5. Sakata, H., Taira, M., Murata, A. & Mine, S. *Cereb. Cortex* **5**, 429–438 (1995).
6. Rigotti, M. et al. *Nature* **497**, 585–590 (2013).
7. Shadlen, M. N. & Newsome, W. T. *J. Neurophysiol.* **86**, 1916–1936 (2001).
8. Andersen, R. A., Essick, G. K. & Siegel, R. M. *Science* **230**, 456–458 (1985).

9. Zipser, D. & Andersen, R. A. *Nature* **331**, 679–684 (1988).
10. Andersen, R. A., Snyder, L. H., Bradley, D. C. & Xing, J. *Annu. Rev. Neurosci.* **20**, 303–330 (1997).
11. Bougou, V. et al. *Cell* **189**, 5304–5317.e7 (2026).
12. Battaglia, P. W., Hamrick, J. B. & Tenenbaum, J. B. *Proc. Natl Acad. Sci. USA* **110**, 18327–18332 (2013).
13. Churchland, M. M. et al. *Nature* **487**, 51–56 (2012).
14. Yin, F. et al. *Commun. Biol*. https://doi.org/10.1038/s42003-026-10369-8 (2026).
15. Wandelt, S. K. et al. *Nat. Hum. Behav.* **8**, 2295–2308 (2024).
16. Kunz, E. M. et al. *Cell* **188**, 4658–4673.e17 (2025).
17. Ibsen, S., Tong, A., Schutt, C., Esener, S. & Shapiro, M. G. *Nat. Commun.* **6**, 8264 (2015).