PERSPECTIVES / **COMPUTATIONAL NEUROSCIENCE**

# Whole-brain, bottom-up neuroscience: The time for it is now

Applying new tools to entire brains, starting with *C. elegans*, offers the opportunity to uncover how molecules work together to generate neural physiology and how neurons work together to generate behavior.

BY **EDWARD BOYDEN, KONRAD KÖRDING**
20 OCTOBER 2025 | 9 MIN READ

**HTTPS://DOI.ORG/10.53053/IPUB1951**
**CITE THIS ARTICLE**

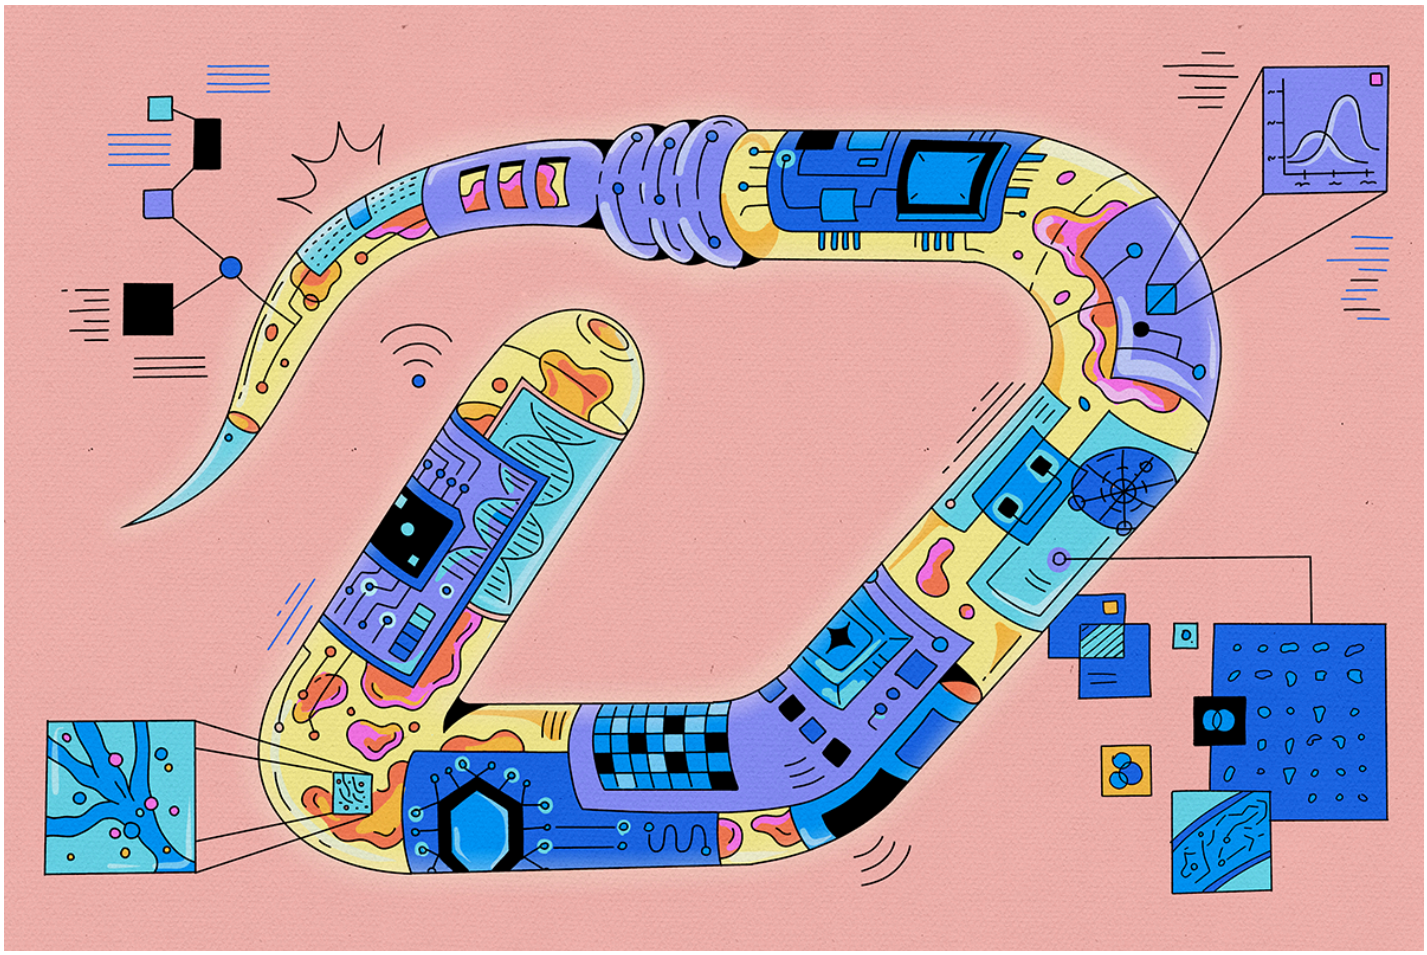

**Building blocks:** Collecting the data needed to create a whole-worm model would require running multiple kinds of experiments on the same animal and acquiring data from many animals.
ILLUSTRATION BY KATHLEEN FU

Take a handful of sand from a beach. Could anyone in 1900 have imagined that this sand could become a microchip powering artificial intelligence that rivals the human intellect? Even the boldest science fiction writers of that era never dreamt of such things. Computing machines, if they were imagined then at all, were crude electromechanical contraptions. Only after quantum mechanics emerged in the early 20th century did we grasp the potential of semiconductors (such as silicon, a major component of sand), which led to the microchip, and thus to modern computing. Once we understood how the fundamental particles of physics interact, innovation exploded: The 20th century saw the emergence of such impactful achievements as lasers, the moon landing and the internet.

When we know the building blocks of a domain and their interactions—or "ground truth," for short—we can design from the bottom up, unleashing untold combinatorial possibilities. In mature sciences, such as physics, we can start with the known properties of components—electrons or steel beams, for example—to simulate and predict the behaviors of systems made out of these components, such as microchips and bridges. Discovery is less about luck and more about design: Start from parts with known interactions; then simulate—and build.

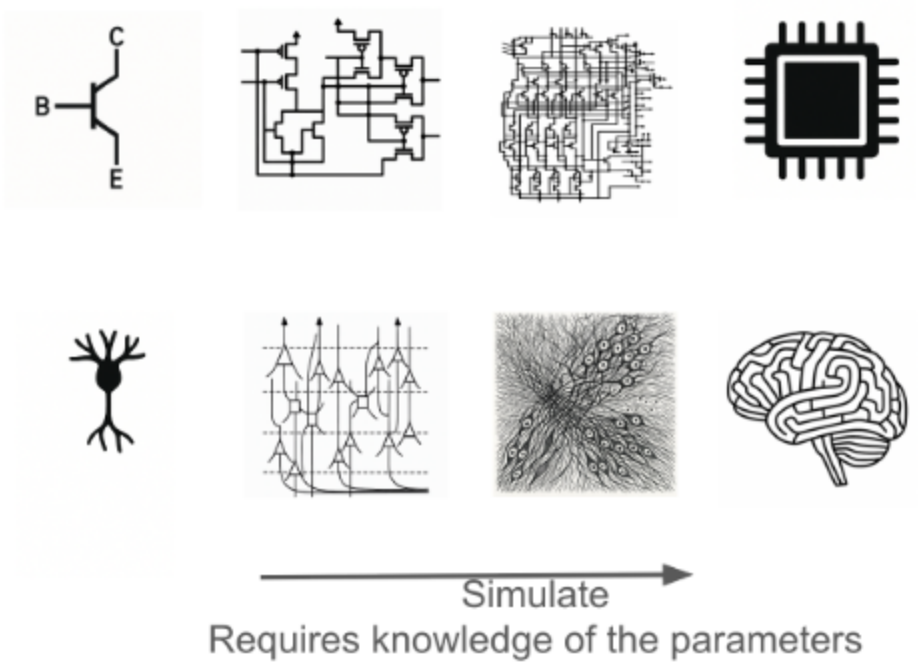


**Bottom up:** In the same way that we can simulate microprocessors from transistors up, we should develop technology to simulate brains from neurons up. Transistors come in few types, whereas neurons have many parameters. Therefore, we need technology that gives us these parameters.

Neuroscience, by contrast, still mostly works top down. We observe an effect and then develop and test a hypothesis. We try to understand feelings, decisions, sensations, movements and even consciousness through this process. Sometimes this works, revealing a mechanism. But the hypothesis space is enormous: hundreds of brain regions, thousands of cell types, perhaps millions of molecule types. In such large spaces, almost all hypotheses are wrong. Testing a hypothesis in isolation rarely leads to understanding how a system works in terms of its mechanisms, and if it does, we only learn about a few of them.

Really understanding how a brain process occurs—or curing a brain disease—remains daunting; there are countless hidden dimensions to confound us. In contrast to the microchip paving the way for the personal computer and then the internet, with each step leading to exponential growth at the next level, major neuroscience discoveries, such as orientation-selective cells or long-term potentiation, did not make the next leap exponentially easier.

Can we take neuroscience to its "microchip moment"? It's difficult to achieve ground truth for brains. Some scientists have succeeded at bottom-up neuroscience in small systems, such as the squid giant axon or the crab stomatogastric ganglion, yielding transformative foundational principles of how brains work. But to understand entire brains, one must not only deal with the heterogeneity of their building blocks but also the vast spatial and temporal scales over which brains operate.


**Fly through:** A *C. elegans* worm has been stained with a dye that binds amines (a reactive group found on proteins, among other biomolecules), shown in gray. Scale bar: 20 micrometers

Here we outline an approach that we think will enable dynamic observation and control of entire brains, as well as the ability to map at the level of molecules, wires and neurons, and we describe how such datasets could yield bottom-up models. To make the argument concrete, and to provide a practical test bed for a first actual attempt, we focus on the worm *Caenorhabditis elegans*. Cracking the worm would help us see how molecular recipes generate neural physiology. And it would enable us to understand how neurons work together to yield emergent brain functions. We also outline how these approaches might

scale to larger brains—including the larval zebrafish, the mouse and even the human—in the years to come.

With its 302 neurons, *C. elegans* feeds, mates, flees and makes complex decisions.

An electron microscopy-derived connectome of the worm has been available for decades. More recently, researchers have developed gene-expression profiles of all the cells of *C. elegans* and recorded and stimulated, with single-cell precision, activity throughout the worm's nervous system. These datasets, though impressive, have limitations. They are too small to use to build whole-worm models that describe the behavior of the worm nervous system from neuron-neuron interactions, and they were collected from different worms, making it difficult to integrate molecular, wiring and activity data.

Collecting the data needed to create a whole-worm model would require running multiple kinds of experiments on the same animal and acquiring data from many animals (**Box 1**). New tools are making this possible. Both brain mapping (mapping all neurons, their synapses and associated molecules) and control technologies (stimulating and recording all neurons) are already operating at the scale of the entire nervous system of *C. elegans*.

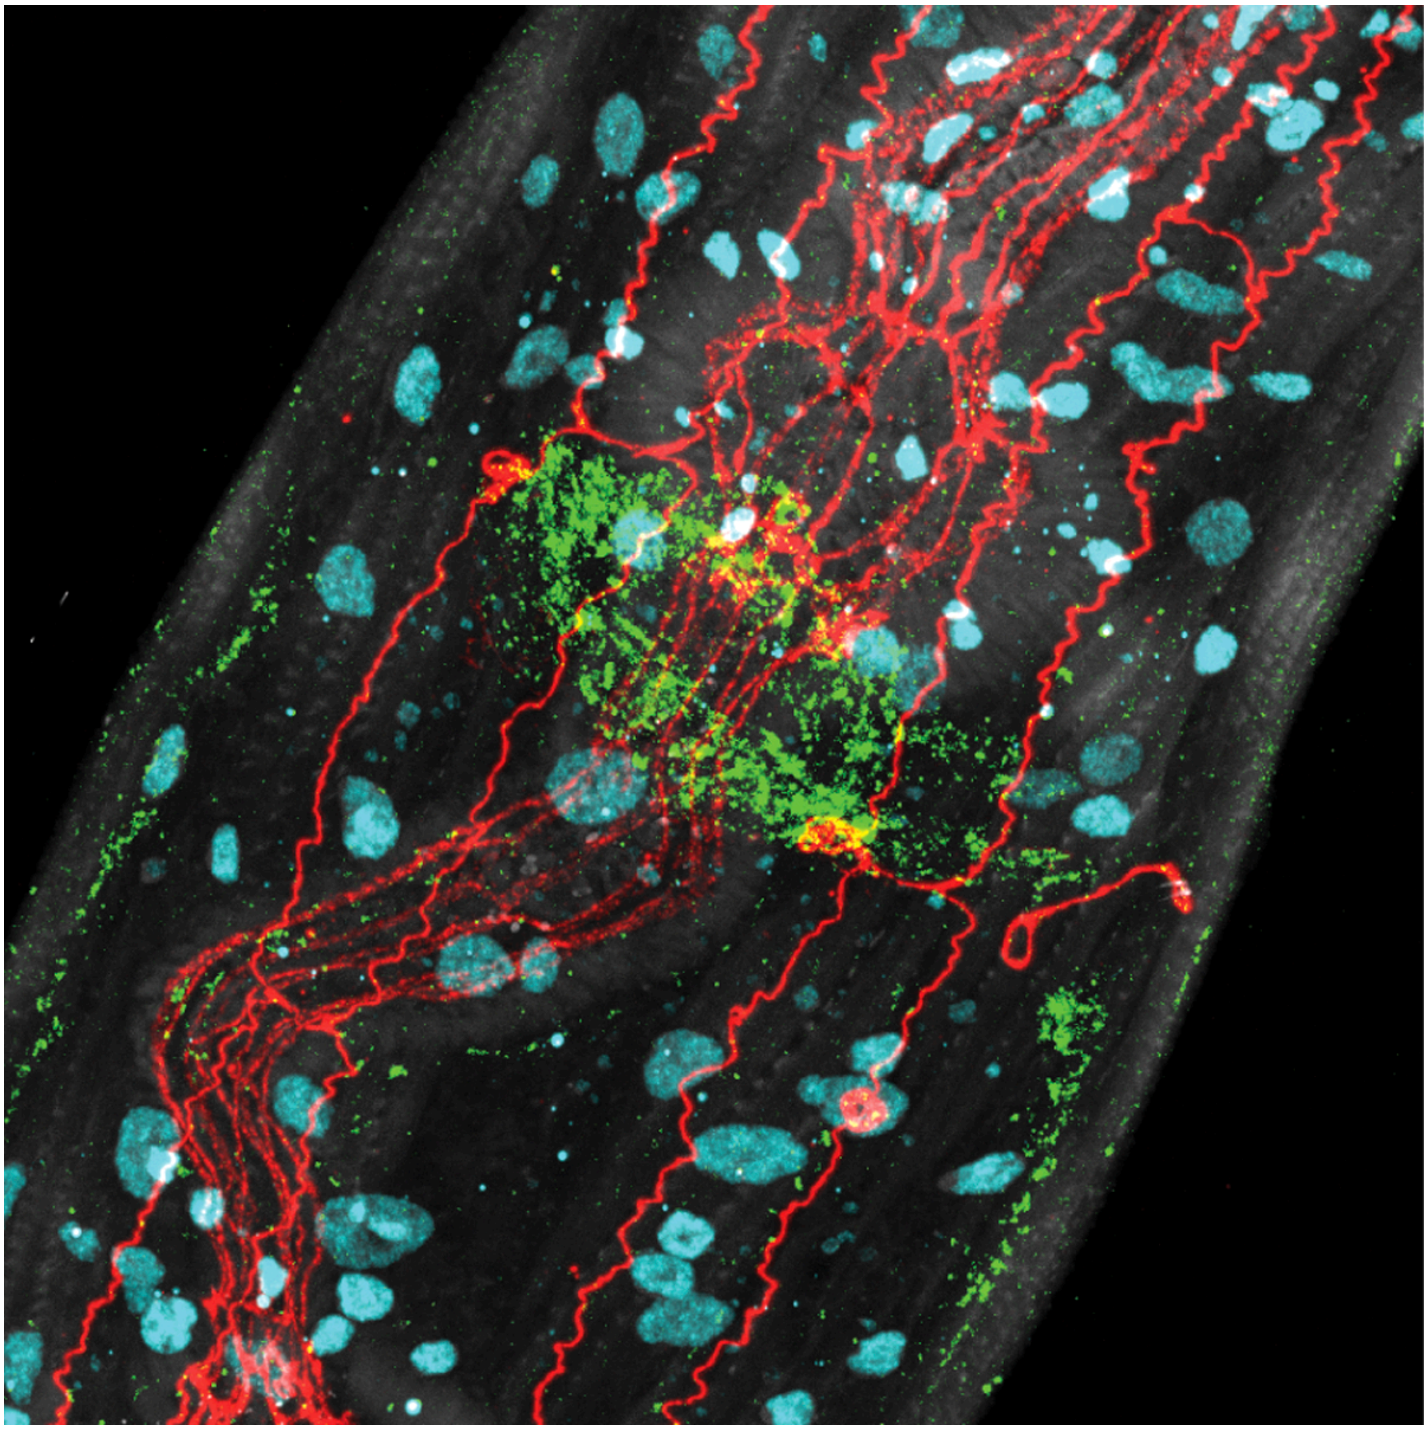

**Expansion microscopy:** This *C. elegans* worm has been expanded in size threefold and treated with stains that target amines (a reactive group found on proteins, among other biomolecules) in gray, a structural protein (DLG-1) in red, a motor protein (DYN-1) in green, and a nuclear protein (DAO-5) in cyan. Scale bar: 10 micrometers (in biological units)

To decipher molecular recipes, researchers could, in principle, apply expansion microscopy pipelines, with machine-learning-facilitated connectomic tracing, protein-shape imaging, multiplexed antibody staining or subcellular spatial transcriptomics to a complete *C. elegans*. How do we learn what each molecule actually does so that we can convert images of *C. elegans* into a simulation? One possibility is to estimate how dynamics within, and interactions between, the mapped molecules give rise to function, using AlphaFold-style or other kinds of computational modeling.

Another strategy would be to study molecular interactions and dynamics through optogenetic protein perturbation and multiplexed live imaging, either in native cells or after reconstitution in heterologous

systems. Of course, there might be many molecules we don't know about yet: Machine learning could help us infer some of these, based on the ones that we do see and perturb. Such datasets could yield bottom-up biophysical models of how the molecules of a neuron give rise to its emergent physiology, via cable equations based on molecular location and kinetics, for example.

To understand how neurons give rise to brain dynamics, we can treat individual cells as building blocks and model their interactions through synaptic and non-synaptic signaling. This could be done by extending the bottom-up models of the previous paragraph into networks of neurons. We could also probe causality by optogenetically perturbing the electrical activity of single neurons (or ensembles) while recording voltage or calcium signals in the rest of the network. When we perturb neuron A, what does it do to neuron B? If we could measure those parameters, we could simulate how the neurons work together in networks.

**Estimating data: How much data do we need to estimate the parameters (ground truth) of a neuron?**
One key way we can get intuition on how much data we need is the approximation that the parameters (**A**) define how the inputs (**x**) linearly affect the output of a neuron y=**Ax**. We then want to estimate A from x and y—a classical *inverse problem.* We can characterize the error $\|\hat{A} - A\|_2$ analytically in such cases:

$$\mathbb{E}\|\hat{A} - A\|_2 \approx \sigma_\varepsilon \sqrt{\frac{p}{n}} \underbrace{\sqrt{\frac{1}{p}\operatorname{tr}(\Sigma^{-1})}}_{\text{conditioning factor } C(\Sigma)}$$

In that case, the amount of data we need for a fixed error norm increases linearly with the number of parameters *p*. It also increases as the data start differing from the inputs being an isotropic Gaussian and become correlated (or more precisely small singular values of correlation matrix become small) as the conditioning factor explodes then. Because real-world data tend to be very correlated, recording from all neurons is probably insufficient to identify even linear weights of a realistic system.

Given the large number of experiments required, and the need for optical control and imaging of every neuron in the nervous system (you can't just divide a living brain into parts and scan them in parallel), *C. elegans*' small brain offers a natural test bed for this kind of work. Indeed, whole-worm calcium imaging and single-neuron optogenetic control during calcium imaging are already in use. Activating individual neurons or sets of neurons and imaging responses throughout the nervous system could generate an input-output map of the worm. Such a map could form the basis for a model in which every input-output function that can measured will act as a piece of the simulation; it would work like a node of an artificial neural network. We can then simulate the whole brain by simulating this identified system. By following such live experiments with the aforementioned molecular and wiring mapping, it would be possible to

infer how molecular recipes lead to emergent physiology. In the beginning, we will need to validate such models by making predictions that we test in the wet lab; for example, we could knock out an ion channel or stimulate a neuron in the model and then compare it to a CRISPR'ed or optogenetically controlled worm, respectively.

Of course, one must test hypotheses derived from the worm in other organisms. Insights from *Aplysia* and the crab stomatogastric ganglion led to general principles—such as how kinases regulate synaptic strength or how circuits achieve homeostasis—that held in mammals. Understanding how an entire nervous system works as an emergent whole will yield predictions that can be tested in mammals, and perhaps even humans.

Over time, as the model makes successful predictions, it might come to be its own resource. Researchers could do drug screens in the model or look at internal states not obviously linked to behavior and model their impact. Students could one day learn neuroscience on a digital *C. elegans*. Circuit motifs distilled from the model might inspire new AI architectures.

And of course, the journey does not end with *C. elegans*. Indeed, integrating these tools in the context of the worm would pave the way for scaling up the mapping and modeling to larger and larger brains—perhaps someday even the human brain—capable of complex decisions, creative thinking and ethical outcomes. Already, efforts have begun to image the voltage of neurons distributed throughout the entire zebrafish brain, to image entire expanded mouse brains and to make living mammalian brains more transparent, which could facilitate more scalable causal perturbation and optical activity recording in live mammals. The journey will be long, thus validating that our approach on the worm would be time well spent.

The microchip was perhaps the most publicly facing outcome of the 20$^{th}$ century reduction of science risk in physics. The microchip enabled college dropouts to found Microsoft, Apple and Facebook. Engineering and market risks remained, but the underlying science was sound. By contrast, curing brain diseases still entails decade-long clinical trials, multibillion-dollar costs and sky-high failure rates. Even then, the "successful" drugs often disappoint. Bottom-up neuroscience—being able to simulate brains from ground-truth parts and interactions—promises to change that reality.

TAGS:
COMPUTATIONAL NEUROSCIENCE, C. ELEGANS, CONNECTOME, MODELING, OPTOGENETICS, TRANSCRIPTOMICS


## Edward Boyden

*Y. Eva Tan Professor in Neurotechnology, MIT McGovern Institute and Yang Tan Collective;*
*Investigator, HHMI*



## Konrad Körding

*Professor of neuroscience, University of Pennsylvania;*
*Co-founder, Neuromatch and the Community for Rigor*


SHARE THIS ARTICLE:

## EXPLORE MORE FROM *THE TRANSMITTER*

CONNECTOME

# Worms help untangle brain structure/function mystery

BY **HOLLY BARKER**
29 AUGUST 2025 | 7 MIN READ

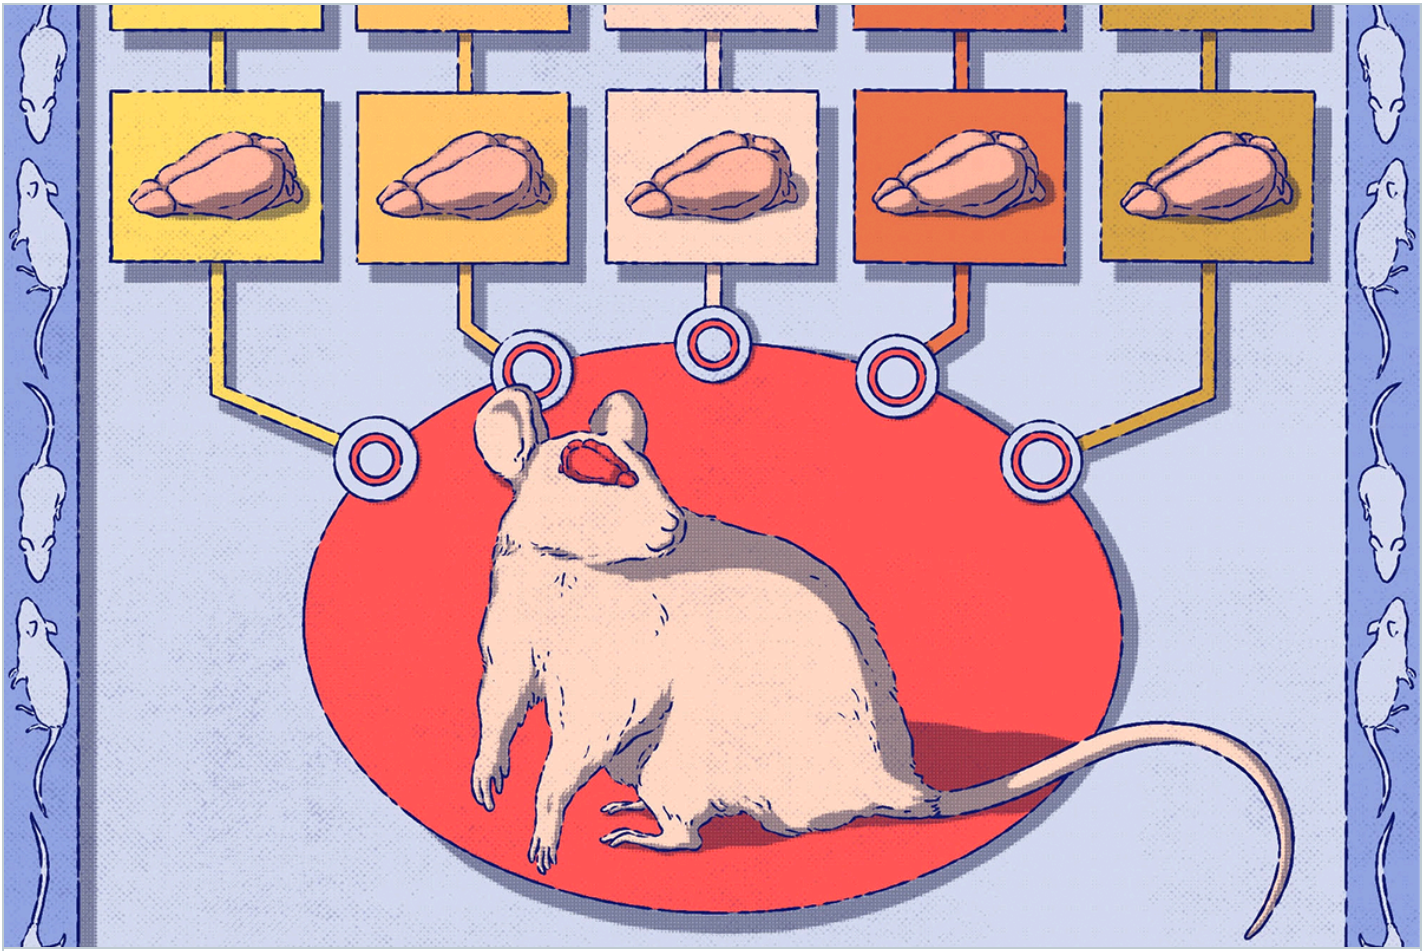

REPRODUCIBILITY

# Reproducibility is a team sport: Lessons from a large-scale collaboration

BY **ANNE CHURCHLAND**
29 SEPTEMBER 2025 | 7 MIN READ

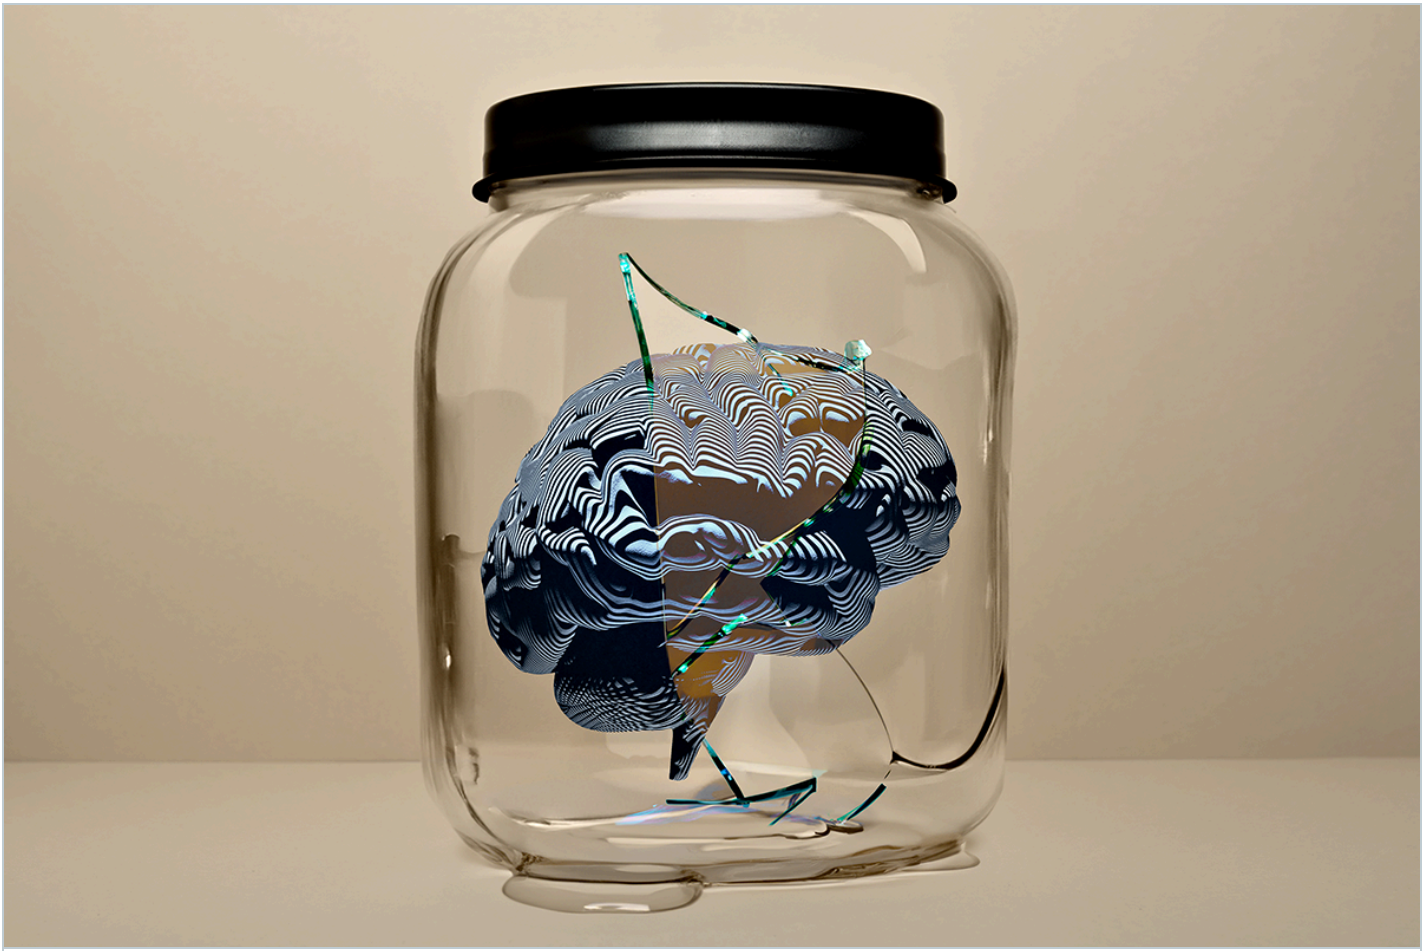

NEUROAI

# Breaking the jar: Why NeuroAI needs embodiment

BY **BING WEN BRUNTON, JOHN TUTHILL**
21 JULY 2025 | 9 MIN READ

## Special report: The state of neuroscience

Where do the field's practitioners think it is headed?

READ MORE

*THE TRANSMITTER* BOOKS

EVENTS

COURSES

SPECTRUM

ABOUT *THE TRANSMITTER*

THE TEAM

PARTNERS AND COLLABORATORS

NEWSLETTERS AND ALERTS

EARLY-CAREER RESEARCHER RESOURCE CENTER

MENTORSHIP DIRECTORY

ANNOUNCEMENTS

CAREERS

RSS

TERMS AND CONDITIONS (SIMONS FOUNDATION)

EDITORIAL POLICIES

AI EDITORIAL POLICY

PRIVACY POLICY

ACCESSIBILITY

CONTACT US

FOLLOW THE TRANSMITTER: