
Thinking about neuroscience
Essays from the field

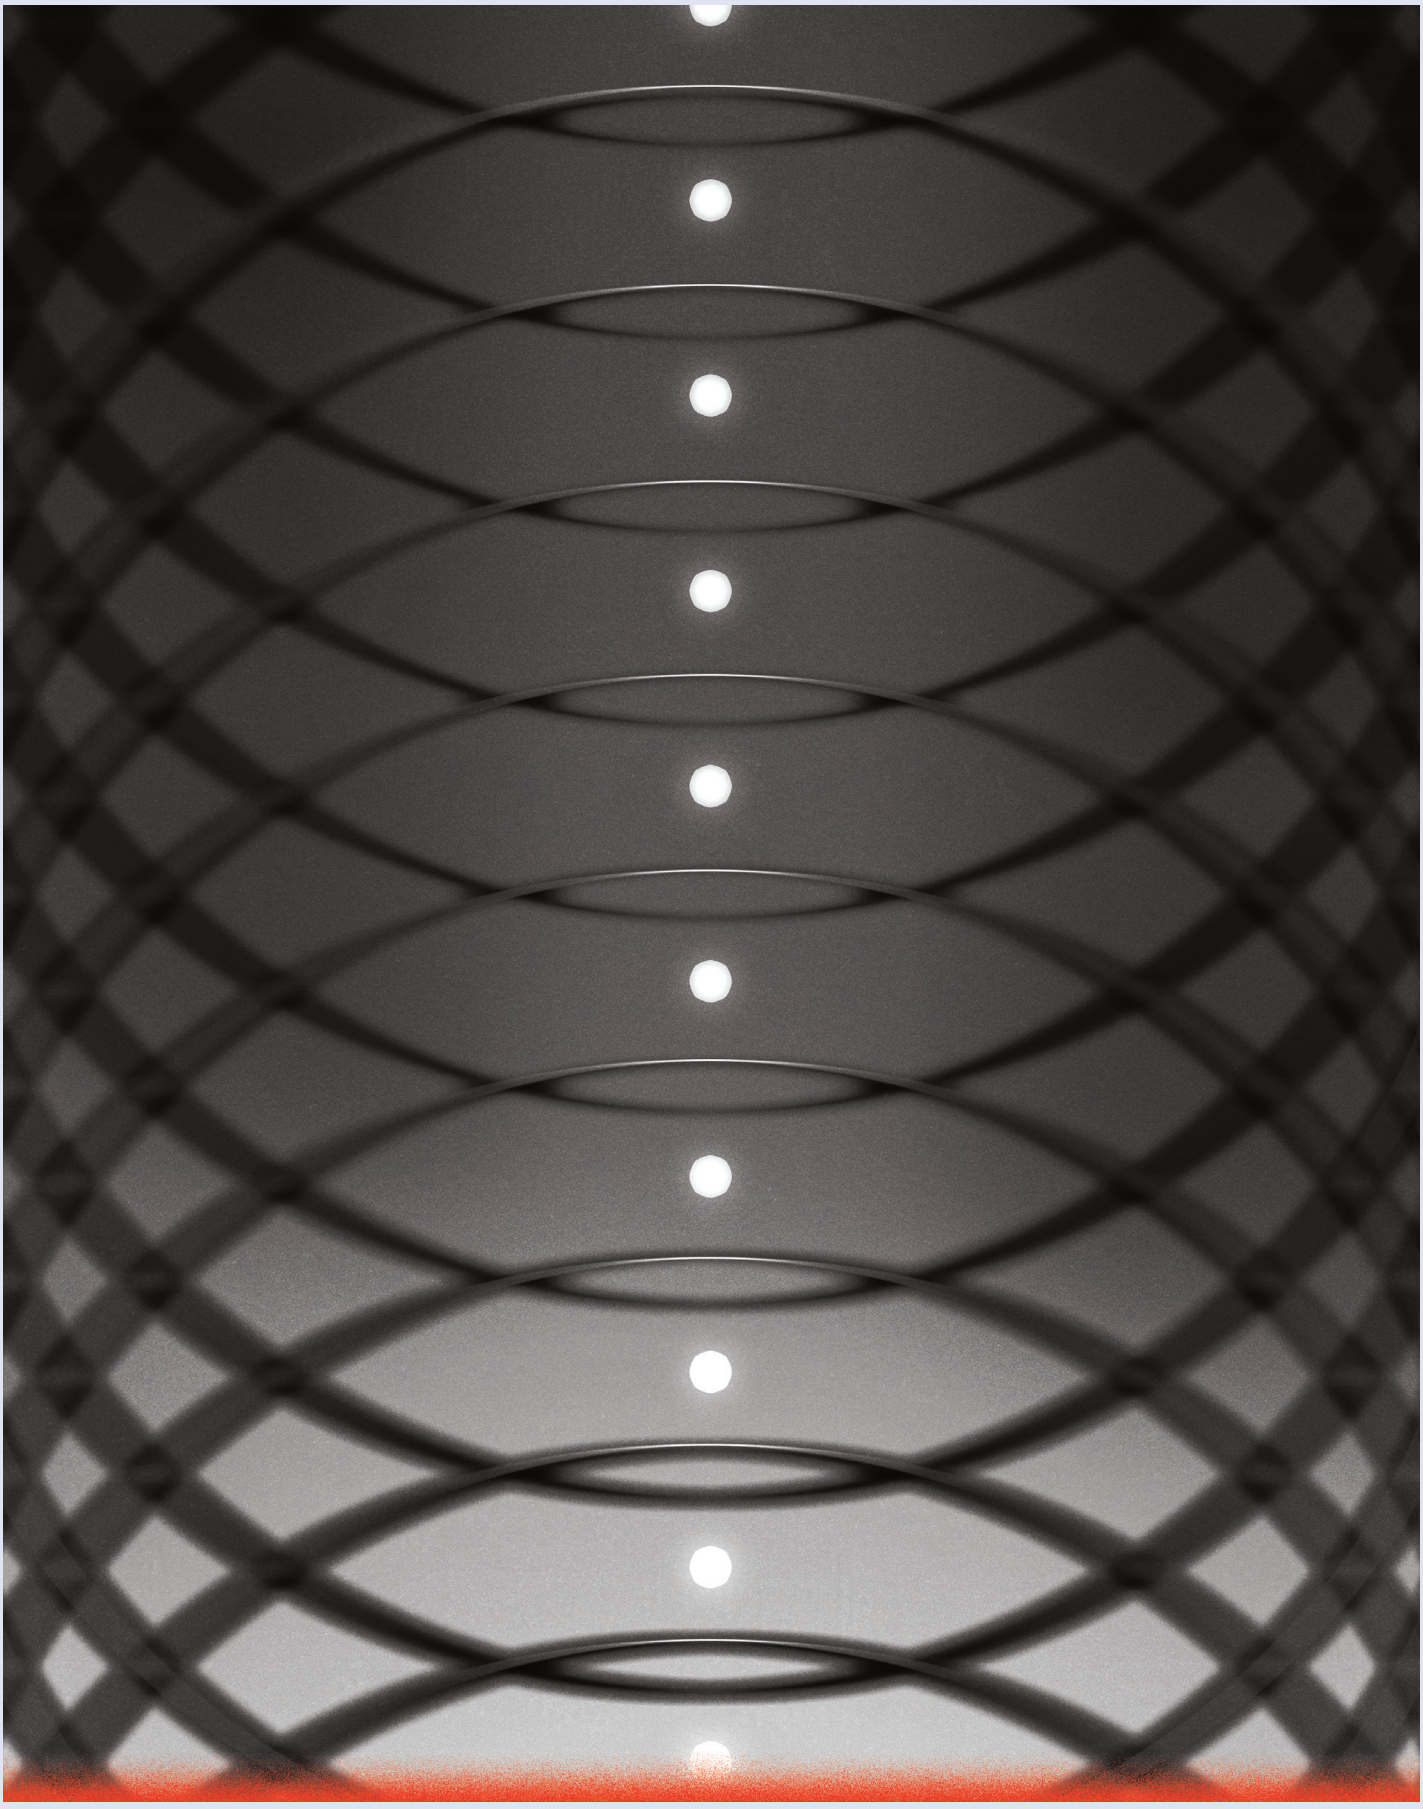

Thinking about neuroscience: Essays from the field

Contents

Transmitter Team

Ivan Oransky	EDITOR-IN-CHIEF
Kristin Ozelli	EXECUTIVE EDITOR
Nicholette Zeliadt	MANAGING EDITOR
Nektaria Pastellas	DIRECTOR, MARKETING AND ENGAGEMENT
Emily Singer	CHIEF OPINION EDITOR
Brady Huggett	FEATURES EDITOR
Rebecca Horne	ART DIRECTOR
Hope Vanderberg	COPY CHIEF
Francisco J. Rivera Rosario	ASSOCIATE EDITOR, OPINION AND COMMUNITY
Angie Voyles Askham	SENIOR REPORTER
Calli McMurray	REPORTER
Shaena Montanari	REPORTER
Sam Schuman	WEB PRODUCER
Jasmine Peralta	ASSOCIATE, MARKETING AND ENGAGEMENT
Iris McNitt	EDITORIAL ASSISTANT
Anya Sahni	ART INTERN

Book Team

Catherine Casalino	BOOK DESIGNER
Ibrahim Rayintakath	ILLUSTRATIONS AND COVER ART

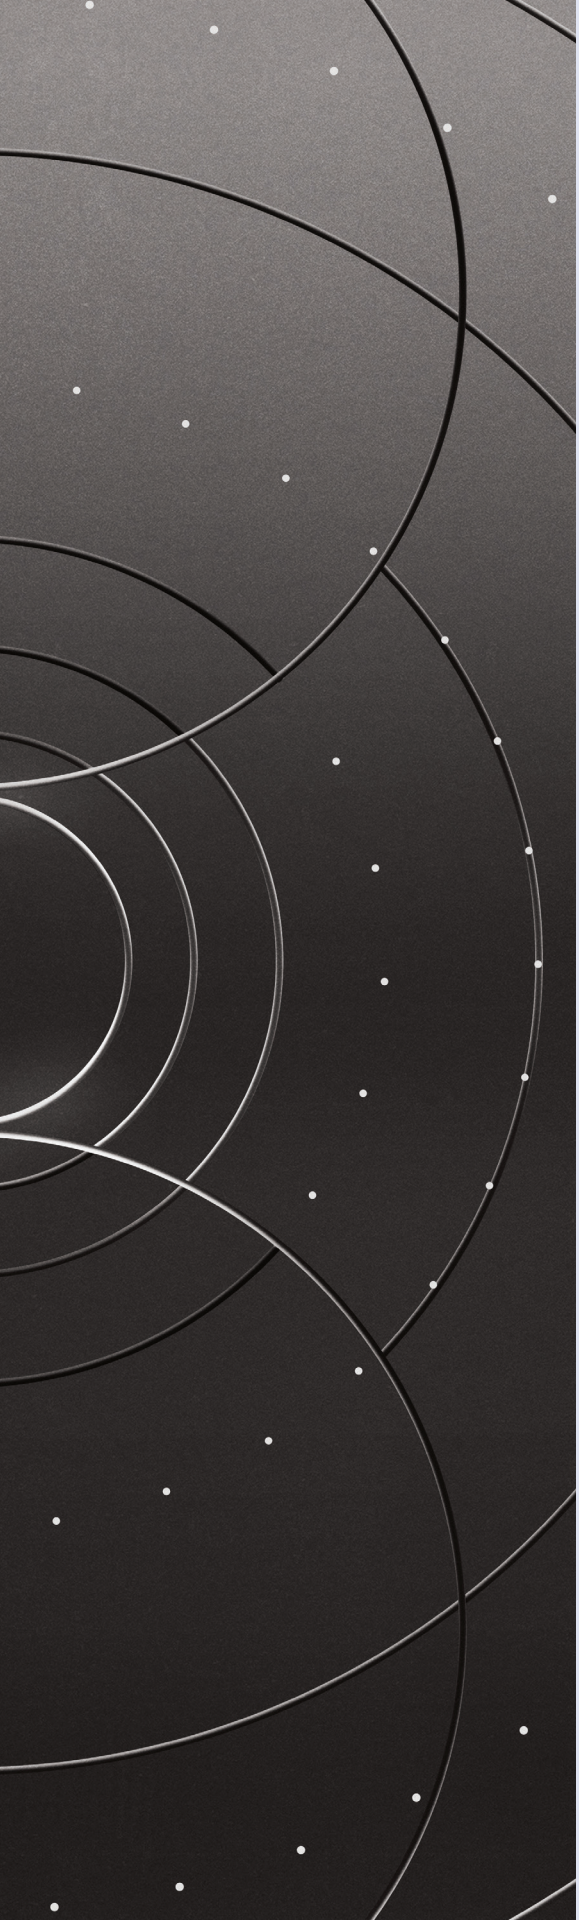

Letter from the editors

Our driving ethos at *The Transmitter* is simple: to be useful to the neuroscience community. To that end, the voices of scientists have been central to our publication—as sources, contributors, columnists and writers—since its launch a little more than a year ago.

For this, our first book, it seemed only appropriate, therefore, to amplify those voices further. In this collection, we feature some of our favorite scientist-written essays from the past year—pieces that together not only capture the breadth of the field and the opinions within it, but also demonstrate the range of essays we run each week.

These pieces offer members of the neuroscience community a place to explore the culture and practice of their field; to debate major questions; to express novel ideas to advance their work; to examine the craft of neuroscience; and to trace its evolution over time.

On page 16, for example, Jakob Voigts implores the field—known for its DIY tendencies—to increase its efficiency by embracing professional help. Megan Peters explains on page 114 how to properly credit long lists of contributors, an increasingly common problem as neuroscience collaborations grow. Anne Churchland and Felicia Davatolhagh outline steps on page 110 to address the continued gender bias in neuroscience citations. And on page 22, Anthony Zador chronicles his efforts to bring together two disparate groups in neuroscience—experimentalists and theorists—at the annual COSYNE meeting.

In this compendium, we also highlight some of the essay series we have developed around critical practical challenges in neuroscience. The Open neuroscience and data-sharing series, for example, delves deep into the hurdles and benefits of adopting more open and standardized practices across a field as varied as neuroscience.

Among our most popular essays are those that tackle big questions in neuroscience, and so we have included several of those as well. On page 42, Nicole Rust analyzes the brain's functional regime, asking whether it functions like a chaotic system. And she asks 14 colleagues—among them Eve Marder and Kanaka Rajan—to weigh in on the implications. On page 122, Francis T. Fallon, Tomás J. Ryan, John W. Krakauer and their collaborators within the Representation: Past, Present and Future group describe the many meanings neuroscientists ascribe to "representations," and advocate for a taxonomy of the term.

It's a lot to take in, which is why we also asked several of our contributing editors to weigh in on where they think the field is headed in the next 10 to 20 years. Flip to page 73, for instance, to learn why Russell Poldrack calls for a "shift from the current focus on data to a heavier focus on theory."

We hope you enjoy this small sample of what neuroscientists have written for *The Transmitter* so far, and visit the site regularly to keep up to date with a view of the field from within.

—The editors

TO READ MORE ESSAYS FROM THE TRANSMITTER:

thetransmitter.org/perspectives/

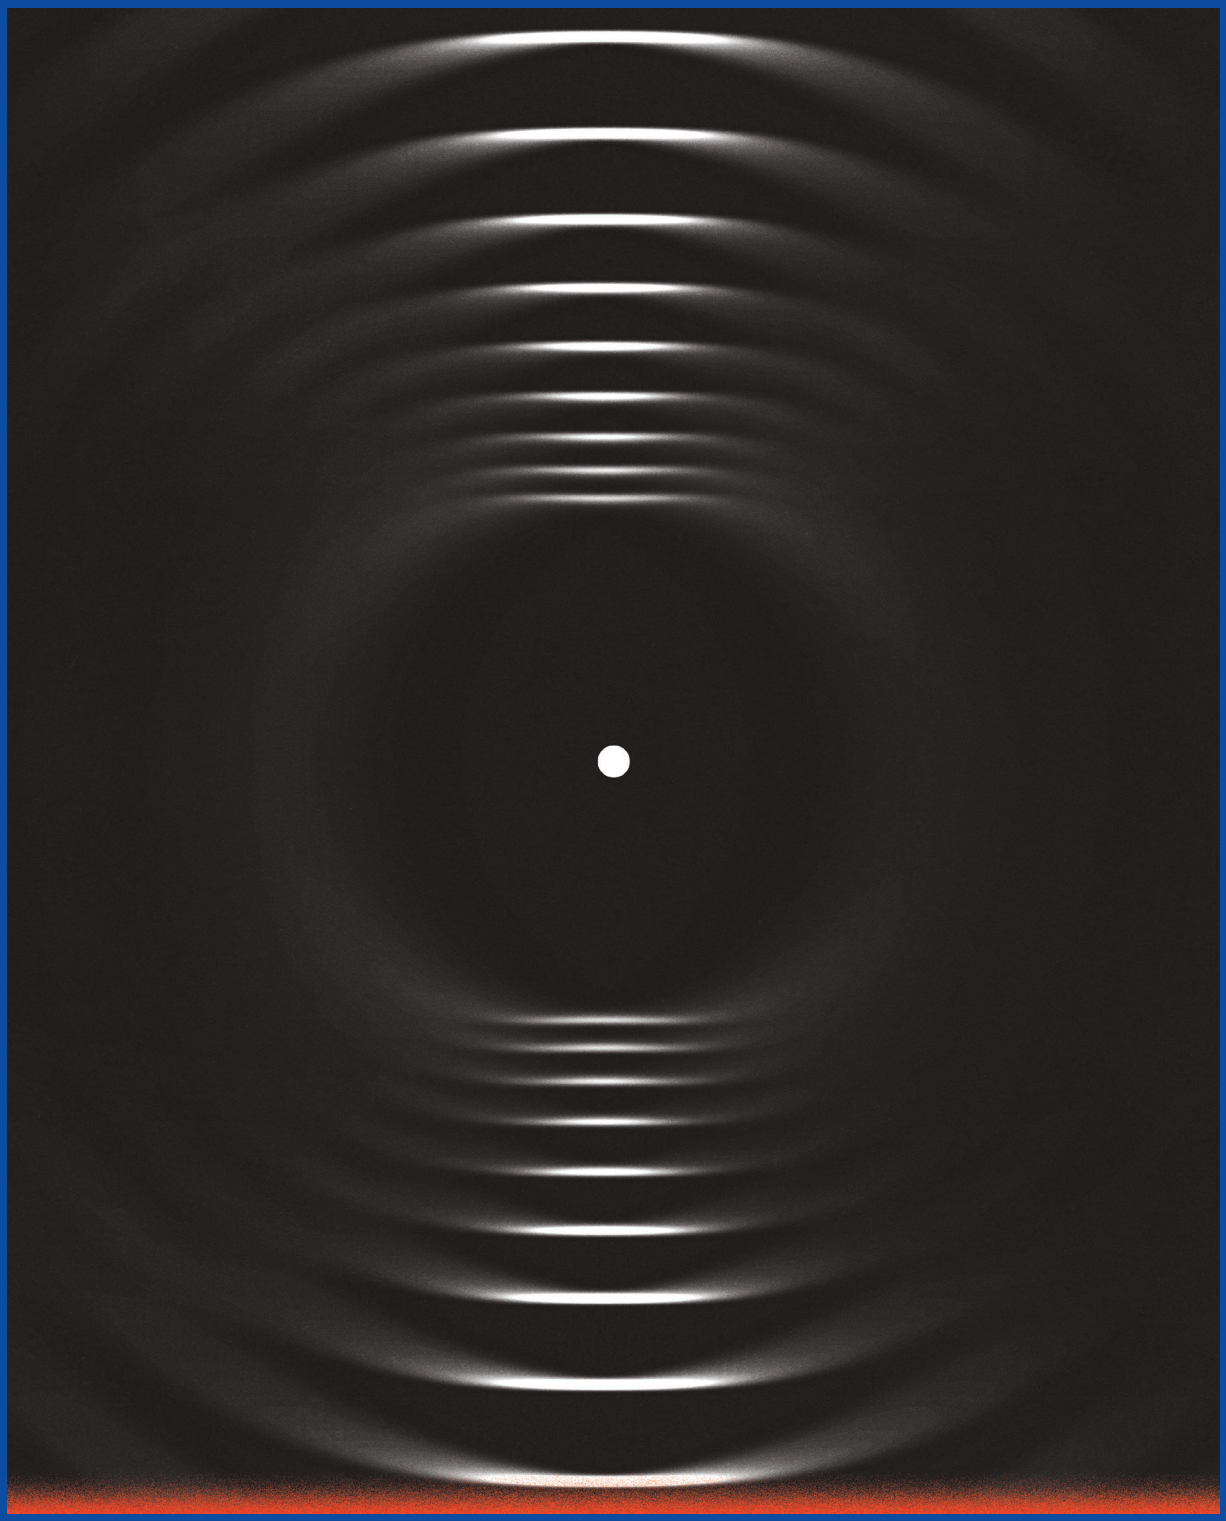

Neuroscience

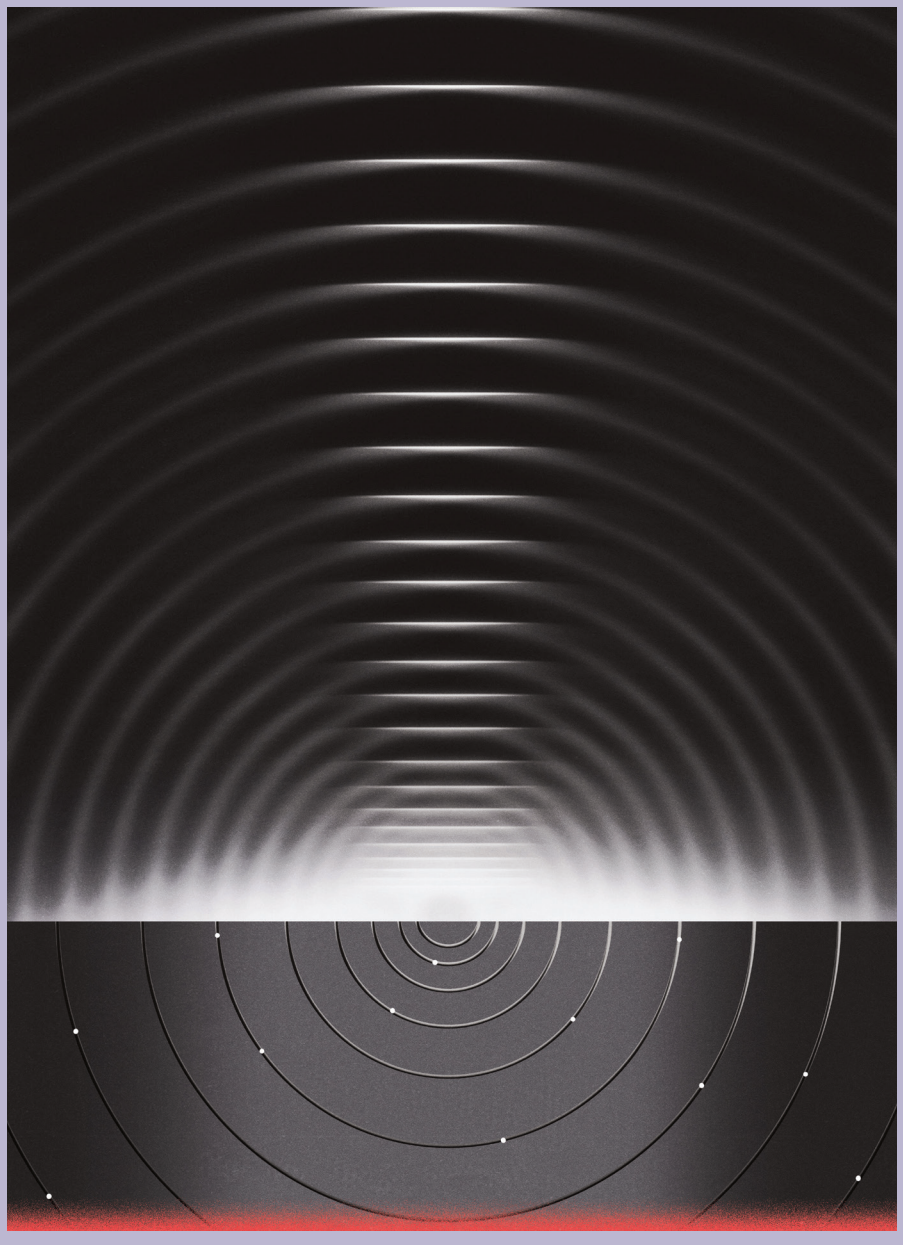

as a profession

Our planet stands on the brink of irreversible change. Neuroscientists need to do something about it.

ILLUSTRATION BY INA JANG

When I launched my new lab at New York University in 2022, I decided to apply my expertise in computer vision to an urgent problem far outside the brain: climate change.

BY GRACE LINDSAY, ASSISTANT PROFESSOR OF PSYCHOLOGY AND DATA SCIENCE, NEW YORK UNIVERSITY

A new professor has to learn many things when starting the job: how to recruit and manage people in their lab, how to organize a curriculum and teach a class, how to navigate departmental bureaucracy and apply to grants. One thing usually not on that list: how to do research in an entirely new field. Yet that is the position I find myself in as a new assistant professor of psychology and data science at New York University.

Although my job application was evaluated mainly on the basis of my work as a computational neuroscientist, I, like an increasing number of neuroscientists I know, have decided to put part of my research efforts toward the dire problem of climate change. To do so is unorthodox in academia and historically would not be well received. But tackling the climate crisis requires radical new approaches.

There is good reason to believe that an influx of new faces could speed progress on some old problems, including the thorny challenges of climate change. While researching for a book I wrote on how mathematics and physics have influenced neuroscience, I repeatedly saw how cross-disciplinary interaction can help expand the space of ideas within a field. In addition to bringing new tools and perspectives, outsiders can also shake up entrenched habits that

could be holding a field back. Through contact with physics and computer science, for example, biologists have come to embrace preprints and open peer-review practices.

The climate work I am interested in takes the form of machine-learning applications for the field of remote sensing, which involves the use of satellites and other devices to identify and study the electromagnetic signature of locations on Earth or other planets. My neuroscience research centers on the visual system, and the models that I build to study it overlap with methods from machine learning and computer vision.

I have been able to use techniques I picked up while studying the brain to analyze satellite imagery for a variety of climate-related purposes. For example, my main project—done in collaboration with the nonprofit Collaborative Earth—focuses on identifying and studying the impact of beaver dams. These structures can play a positive role in adapting to climate change through their effects on floods, wildfire and soil.

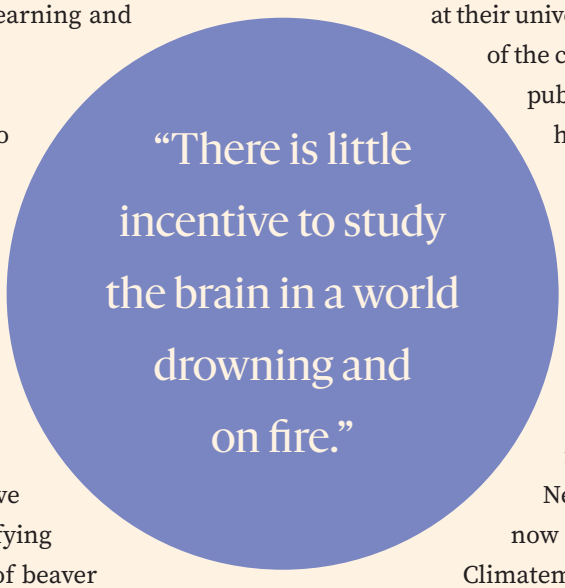

I am not the only one to take this turn. Other neuroscientists have found their own creative paths into climate work. Cognitive neuroscientist Adam Aron turned away from his work on response control to focus on the psychology of climate-change activism; he now runs the Climate Psychology and Action Lab at the University of California, San Diego, and has published a book on the climate crisis. Jeremy Freeman was a group leader at the Janelia Research Campus and director of computational biology at the Chan Zuckerberg Initiative before leaving to start CarbonPlan, a nonprofit organization that provides scientific analysis of climate policies and proposals. Anne Urai, assistant professor of neuroscience at Leiden University in the Netherlands, who studies sensory decision-making, is teaching neuroscientists what they can do in their labs and at their universities to tackle all aspects of the climate crisis; in addition to publishing on the topic, she has organized a Slack group for brain scientists interested in climate action.

Several neuroscientists (myself included) who helped build the online computational neuroscience summer school Neuromatch Academy are now also helping to develop Climatematch Academy, a school that follows the same format as Neuromatch but teaches concepts, coding and research topics related to computational climate science.

What has driven us, as scientists who have thus far dedicated our careers to the study of the brain, to pivot into climate work? For me, I can say that fear and anxiety were certainly motivators. The data are hard to ignore as a scientist, and the data on the climate crisis paint a grim but clear picture of our collective future, even with

efforts to curb the effects underway. The societal and material change still needed to counter climate change is enormous, and society changes only if the people in it do.

Academics, as people who can understand the science behind climate change and have the skills to act (and through teaching also galvanize others to act) reasonably feel a sense of duty. To continue on as though the climate crisis is not occurring would be, in a way, dishonest. It also doesn't hurt that academics usually like to learn, and diving headfirst into an area as big and broad as climate change can be exciting, even if scary.

As I have found, integrating into a new research community presents obstacles: What are the standards and best practices here? Where are the most pressing research gaps? Will my work be taken seriously? Being a newcomer is not easy, but I have felt overwhelmingly welcomed by the community of people working on climate change. Everyone involved knows it is an all-hands-on-deck situation, and they are frequently pleased—relieved, even—to see people from all walks of life take proper notice of the looming problem in front of them. Being upfront about my background and motivations for entering the field has also resulted in people guiding me toward the right resources and collaborators. There is no sense of imposter syndrome when you aren't trying to be an expert but just trying to help.

Academia on the whole, however, is not known for dynamism and adaptability, certainly not on the time scale needed to tackle the climate crisis. Mid-career transitions, labs with two unrelated lines of research, a focus on incremental applied work—all of these possibilities are needed to enable more academics to help fight climate change, even though none are advisable from a careerist perspective. How will I handle running a lab in which some students are studying attention in the visual system and others are studying aerial imagery? It is a problem with little precedent, and I will need to sort it out. How academia can make it easier for someone like me to try to tackle these problems is something we all collectively need to sort out.

I am optimistic that we as a community can make this work. Scientists are judged on all things—papers, grants, promotions—by other scientists; we therefore have the power to set our own priorities. Climate change is an urgent, global problem, and we need to normalize discussions and action around climate change in all areas of science and academia. There is little incentive to study the brain in a world drowning and on fire. I feel privileged to have had any role in helping carry the torch of scientific progress thus far. To keep it burning for centuries more, we—every single one of us—need to do more than just make progress in our field. We need to act to preserve a healthy and sustainable planet on which future scientists can thrive.

Why (and how) we need to professionalize neuroscience

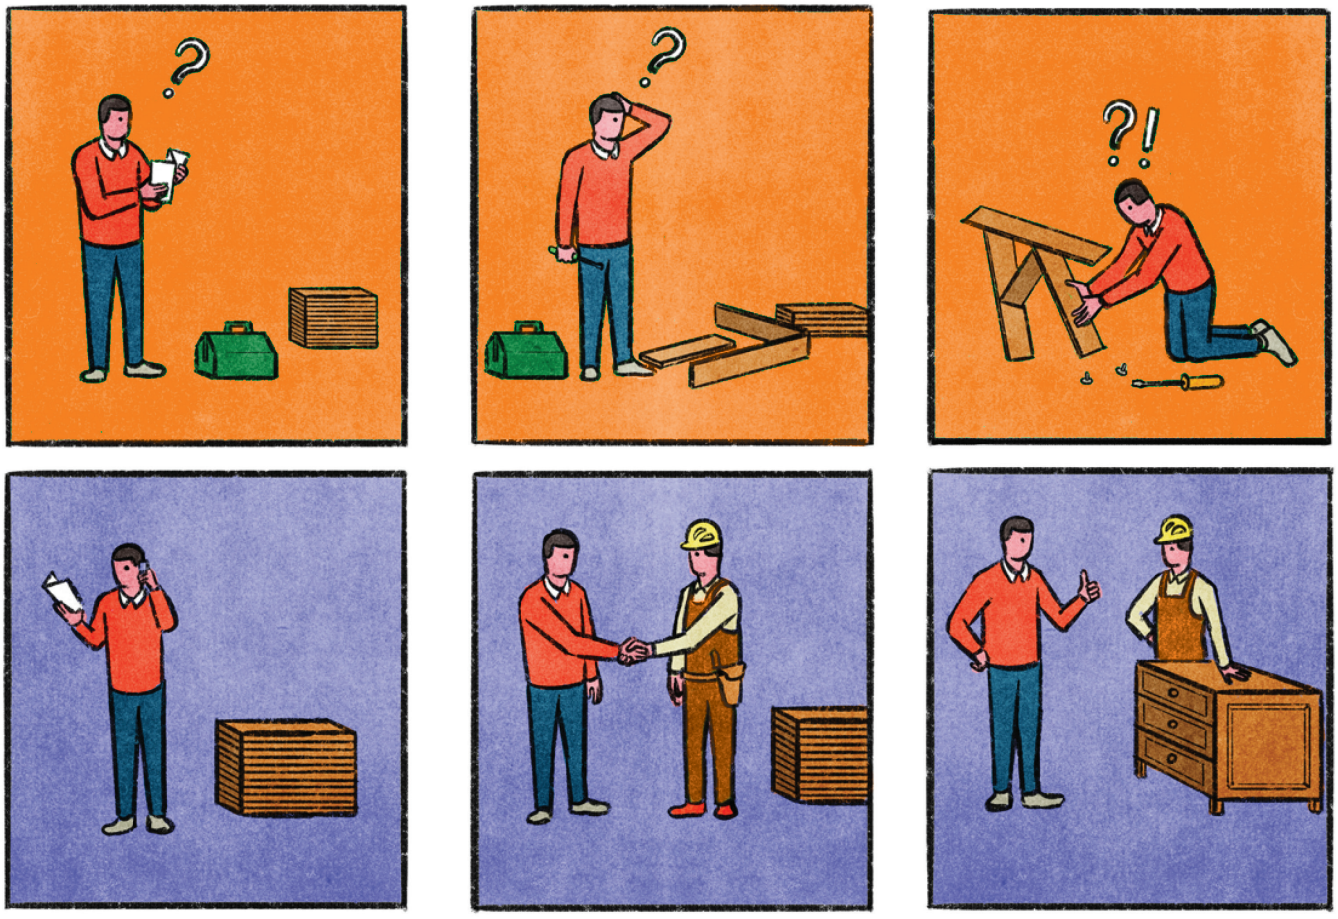

ILLUSTRATION BY KLAUS KREMMERZ

Moving away from the field's do-it-yourself ethos and embracing professional technical expertise will make research more efficient.

In our daily lives, we are accustomed to paying experts for help. Most people recognize it's much faster and safer to hire mechanics to repair our cars than to attempt to do it ourselves. But the same is not true for today's neuroscience, in which a lot of technically demanding work is performed by non-experts.

The resulting inefficiencies are obvious to anyone who has spent time in a lab. Analysis pipelines are hastily written for one paper and then abandoned and redeveloped from scratch by the next postdoctoral researcher. Experiments are performed by research associates after minimal training, with high failure rates seen as an unavoidable part of the process. Multi-lab data infrastructure is built by researchers with no formal IT training and with no plan for supporting it over the long term. This model enables small teams to carry out multidisciplinary projects. But it carries large—and often unseen—costs in funding, time and opportunity, and it ultimately leads to less robust scientific results.

To remedy these costs, we need to enable a culture of professional expertise—a means to hire experts to advise us on, or to carry out, some of the technical tasks that are needed in modern neuroscience. This will require a cultural shift in the field, including new types of grant support, new

BY JAKOB VOIGTS,
GROUP LEADER, HOWARD HUGHES MEDICAL INSTITUTE

career paths for technical expertise and changes in how publications assign credit.

Over the past few decades, neuroscience has become so broad and technically sophisticated that individual researchers can no longer fully understand the technical foundations of their experiments. The average systems neuroscience project, for example, requires in-depth knowledge of animal surgery; mechanical, optical and electrical engineering; statistics and computer science. Neuroscientists very rarely have deep proficiency in all of these domains. Research involving custom tools and procedures, such as surgeries and electrophysiology, as well as novel methods, including viruses, algorithms and custom-designed microscopes, can be especially challenging. Unlike with large commercial devices, such as MRI scanners, appropriate use of these tools requires in-depth knowledge of many technical details.

We underestimate the price of the suboptimal experiments that result from this lack of expertise. Small mistakes can have large consequences: The wrong grounding scheme on an electrophysiology implant can mask reward responses. Improperly soldered connectors can lead to missing weeks of data. A math error in a laser controller can destroy months' worth of samples. In addition, fear of such mistakes incentivizes researchers to stick to methods they know well, which slows innovation. Even more seriously, applying invalid statistical methods; forgetting to include key controls; or failing to correct systematic artifacts produced by the wrong virus serotype, buffer, microscope drift or inappropriate data pre-processing can lead to incorrect results and wrong scientific inferences.

For all these examples, there are experts who would be able to spot and to fix the problems if only they were involved in the projects, be it at the start of a project to help plan, when encountering an issue, or to outsource specific tasks.

Neuroscientists have been reluctant to relinquish their DIY ethos for a variety of reasons. One is the belief that tinkering with methods and pushing techniques beyond known limits or intended purposes is an important method of discovery. Learning about the technical details of our work, and even learning from mistakes, can make us better scientists, particularly for students and postdocs who are explicitly expected to pick up new skills. But not all technical training makes one a better scientist. The most useful types of in-depth technical training are different for everyone. Training electrophysiologists to understand electrical engineering concepts is useful. Expecting them to also become mechanical engineers will not add to their ability to reason about the input impedance of neurons but will instead lead to wasted months trying to fix an incorrectly homebuilt recording rig.

Second, expert help can seem expensive. But it's often cheaper in the long run than having a trainee spend months solving a simple problem, such as a faulty injector or the wrong glue. Beyond the wasted resources—salary, facility costs, supplies—the opportunity cost, missed deadlines and compounded career implications of such delays are bigger still. Economically speaking, avoiding such delays should then be worth a lot of money.

"Expert help can seem expensive. But it's often cheaper in the long run than having a trainee spend months solving a simple problem . . ."

Scientists also often overestimate their ability to become and remain experts across too many domains, and underestimate the amount of time required to solve technical problems. This results in a belief that asking a colleague with relevant expertise for help is enough, when in fact many problems require multi-day visits to diagnose issues, write software or train lab staff.

Putting the reluctance of individual scientists aside, adopting a culture of expertise will also require some shifts in the field. The current funding and publishing system punishes specialization and undervalues technical expertise. For example, most published papers have one (or at best a few) first and last authors. As long as we hang on to the idea of singular intellectual ownership as the main currency in neuroscience, people are incentivized to shoulder as much of their own project as they can rather than spending significant time helping someone else's. A more granular means for giving credit and attribution—one that acknowledges that neuroscience is a team sport—would improve scientific progress.

The field also needs to expand access to expert help. In some domains, institutes such as the Howard Hughes Medical Institute's Janelia Research Campus in Ashburn, Virginia, where I work; the Allen Institute in Seattle, Washington; and the Sainsbury Wellcome Centre in London, England, demonstrate the power of in-house technical expertise. But at most universities, core facilities serve far too few labs to specialize in narrow areas and are permanently either over- or undersubscribed, which often leads to them getting shut down as too costly for the work they provide. This problem would be solved by opening them up to external work, allowing more specialization and evening out the workload, as well as charging sustainable rates for their work, effectively turning them into companies.

Currently, the main source of outsourced technical expertise in neuroscience is tool manufacturers—they provide support and, in some cases, training to scientists. But this is usually tied to purchases, and such companies are often unclear on how to engage with technical work beyond their own tools. To remedy this, we need more flexibility, by making consulting

contracts for bespoke, project-specific services more common. Technical consulting is already commonplace in some areas, such as for server administration, vector cloning, mouse transgenics or air table installation. In systems neuroscience, some small, domain-specific consulting companies already offer such services, including Open Ephys and Aquineuro for freely moving imaging and electrophysiology, and Independent NeuroScience Services for microscope design.

The evolution of Open Ephys, of which I am co-founder, reflects the need for these kinds of services. The organization launched in 2010 as a nonprofit that disseminated scientist-built, open-source devices and software for systems neuroscience, with the aim of reducing the amount of time trainees had to spend getting their experimental rigs up and running. It has recently expanded its offerings, beginning to provide consulting services ranging from custom design work to lab-specific training modules. (I receive no financial compensation for my role at Open Ephys.)

These companies show that this general approach is viable, but we need many more. For the majority of technical issues, paying for expertise is not yet an option.

To expand the market, funding agencies will need to include a means of paying for expertise. Currently, grants typically include funds only for equipment and scientist salaries. If they allow consulting fees, budgets are often tailored to student and postdoc salary levels. Professional technical consulting will be more costly—people with the relevant expertise need career paths that offer stability, salaries and work environments that are competitive with industry. However, these costs would be offset by removing needless delays, reducing the cost of the project overall and increasing the quality and robustness of scientific results.

In sum, a cultural shift that increases the role of technically demanding scientific work as a career path—in academic labs, companies or publicly funded organizations, or as consultants—would be good for science, good for trainees and good for funding agencies seeking to increase the impact of each dollar they grant.

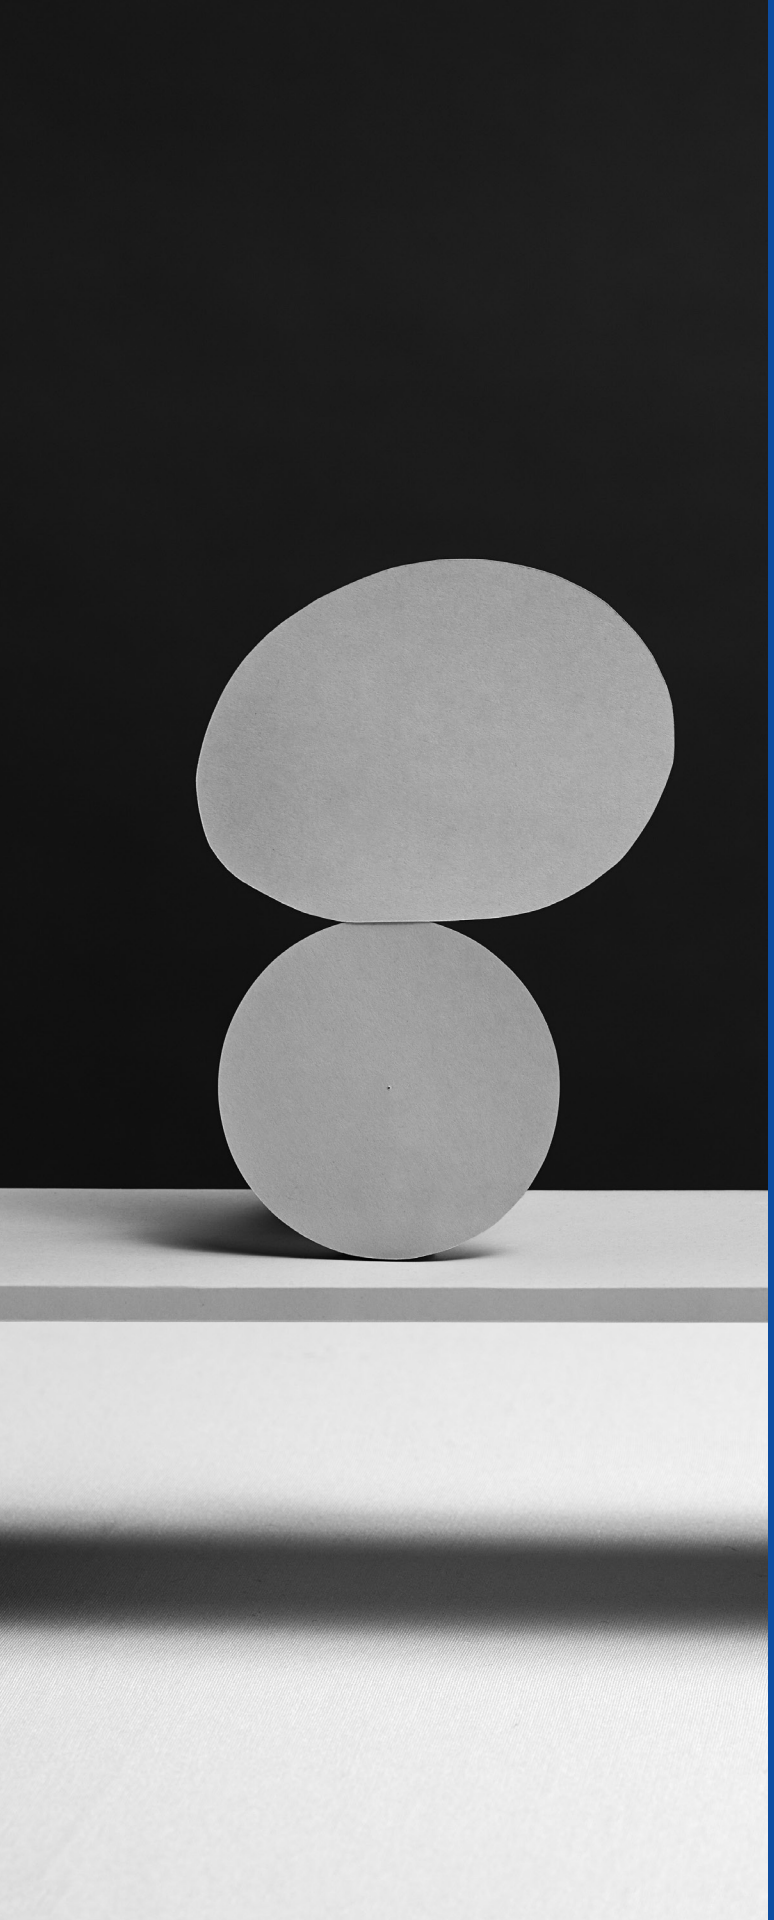

How to explore your scientific values and develop a vision for your field

thetransmitter.org/craft-and-careers/how-to-explore-your-scientific-values-and-develop-a-vision-for-your-field/

As a new professor, I was caught off guard by one part of the job: my role as an evaluator.

BY GRACE LINDSAY, ASSISTANT PROFESSOR OF PSYCHOLOGY AND DATA SCIENCE, NEW YORK UNIVERSITY

The origins of COSYNE: Building a community

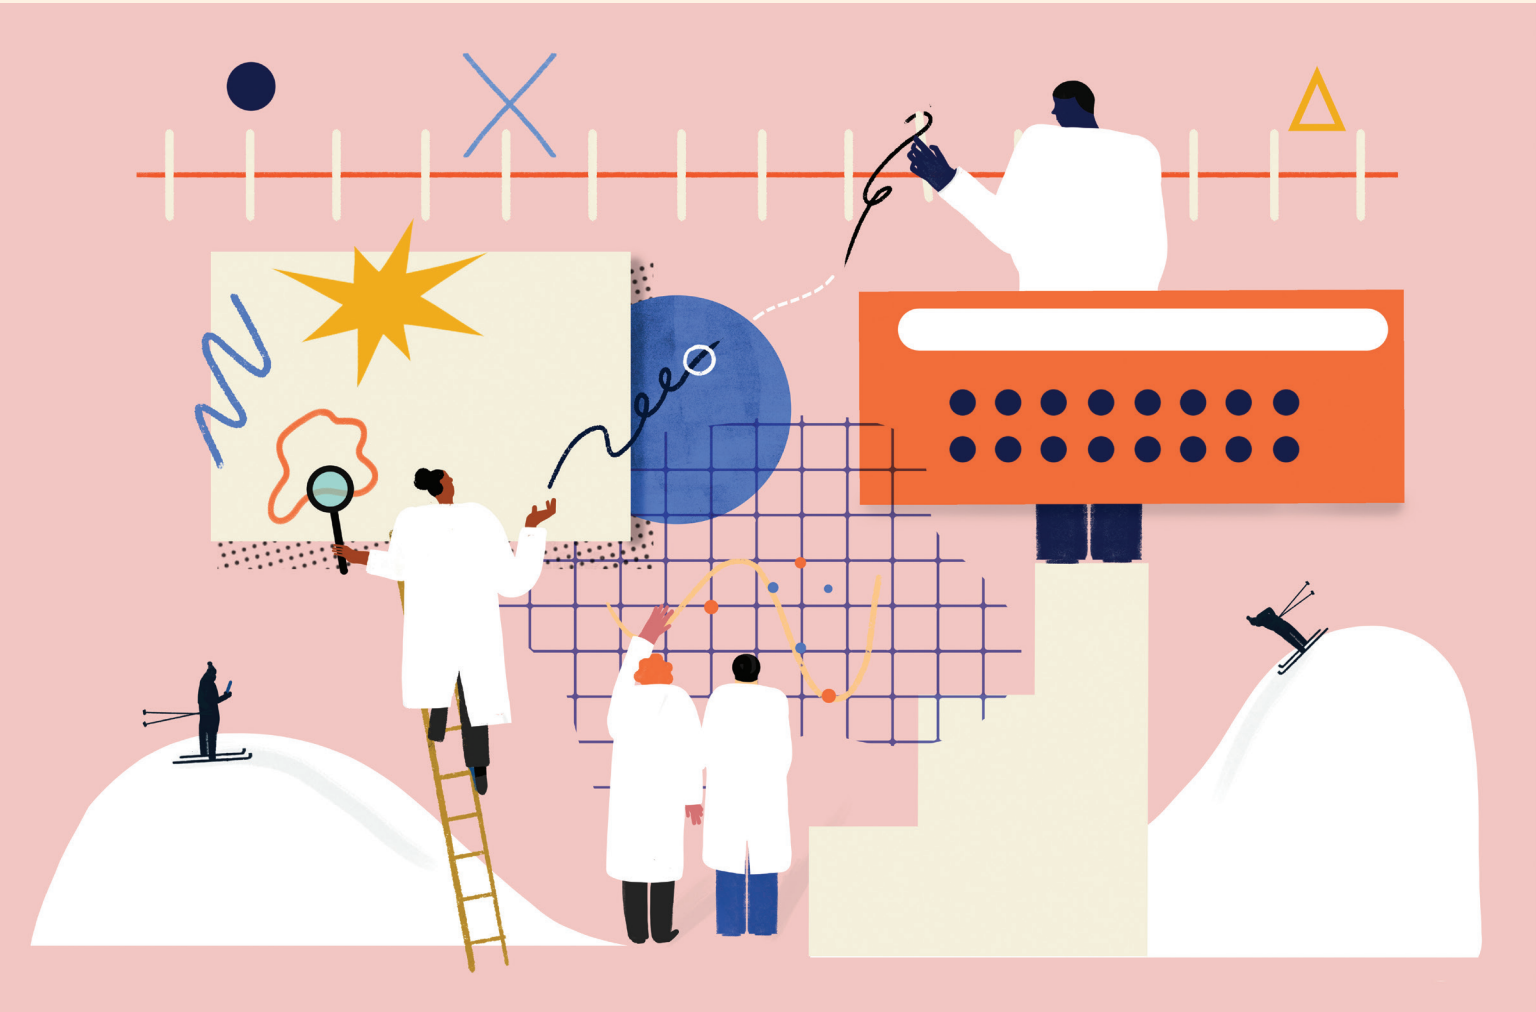

ILLUSTRATION BY MICHELA BUTTIGNOL

Thirty years ago, theoretical and experimental neuroscientists rarely went to the same conferences. So I helped launch a meeting to get them talking.

BY ANTHONY ZADOR,
PROFESSOR OF BIOLOGY, COLD SPRING HARBOR LABORATORY;
CONTRIBUTING EDITOR,
THE TRANSMITTER

Just like you never forget your first love, you never forget your first conference. For me, it was the invitation-only Snowbird Meeting on Neural Networks in 1988, held at the premier ski resort of the same name. Snowbird was scientifically intense, with many of the most engaging discussions continuing on the chair lifts. It brought together a core group of researchers—mostly trained in physics, computer science and engineering, including Yann LeCun, Geoffrey Hinton, Sara Solla, Terrence Sejnowski and John Hopfield—all thinking about how to compute with networks of simple processing elements (artificial neural networks). And it defined for me what a scientific meeting should be like.

Sixteen years later, that early experience inspired me to help found a new meeting, called Computational and Systems Neuroscience (COSYNE), which this week celebrates its 20th anniversary.

The story of COSYNE begins in the late 1980s. At the time, computational neuroscience was still coalescing as a community. Experimentalists were skeptical of theory, and computational neuroscientists often found themselves in a defensive position, repeatedly compelled to debate questions such as: "What good is theory in neuroscience

anyway?" One of the few meetings at which computational neuroscientists felt welcome was the Conference and Workshop on Neural Information Processing Systems (NIPS, later NeurIPS). That meeting evolved from Snowbird as an interdisciplinary gathering of researchers studying both biological and artificial neural networks, but, unlike Snowbird, it was open to all.

NeurIPS was my mainstay meeting during my graduate-school years. But by the mid-1990s, it had shifted its focus to pure machine learning, with neuroscience-related work increasingly treated as an afterthought. As a postdoctoral researcher in the theory-friendly experimental lab of Chuck Stevens, I missed the cross-fertilization of ideas I had found at NeurIPS. So, without a conference to call home base, I started planning a new meeting—Neural Information and Coding (NIC)— to welcome researchers at the intersection of neural theory and experiment.

That was easier said than done. Beyond being simply a place to share information, conferences are essential for creating and nurturing communities. Most people grow up in a single scientific community and don't understand, in practice, how individual communities differ—the structures that establish and reinforce them, the assumptions on which they are built. Language is often a major force in both shaping communities and erecting their borders. A molecular biologist talking to peers, for example, can use the term "restriction enzyme" (a protein that cuts DNA at a specific sequence) without needing to define it, just as a computer scientist can refer to an algorithm whose temporal scaling is "O(N2)" (indicating that the amount of time needed to run the algorithm increases as the square of the number of inputs); but each would likely need to explain those terms when trying to communicate with the other.

My goal in launching NIC was to bring together two communities: theorists and experimentalists. To do that effectively, I realized, they would need to begin to learn each other's languages.

I modeled the first NIC, held in 1996 in Jackson Hole, Wyoming, after that memorable Snowbird meeting—small, invitation-only and with excellent skiing. I thought carefully about how to tackle the language issue and how to engage and excite both groups. Equations, in particular, can form an impenetrable communication barrier for people without the requisite mathematical background. As legend has it, Stephen Hawking was famously advised that each equation he included in his bestseller "A Brief History of Time" would cut his sales in half.

Suspecting that theorists might be more tolerant of jargony talks than experimentalists would be, particularly because theoretical "jargon" often takes the form of equations, I skewed the oral program toward experimentalists. That bias—and likely the location—helped make the meeting a success, and that success continued during subsequent meetings over the next few years in Big Sky (Montana), Snowbird (Utah), Grindelwald (Switzerland) and Les Houches (France). (Lift tickets were much more affordable in those days—less than $20 per day in 2024 dollars—even for the struggling postdocs who organized and attended these meetings.)

Several others joined in as organizers, including, notably, Alexandre Pouget, who would also go on to play a central role in COSYNE. By the early 2000s, word of NIC had spread, and there was growing interest in founding an open, rather than invitation-only, version. In 2004, we held the first annual COSYNE meeting at Cold Spring Harbor Laboratory with 200 participants. Modeling ourselves after NeurIPS, in 2005 we adopted a two-part format: a single-track meeting in the city and parallel workshops at a ski resort.

In COSYNE's early years, a small group of us handled all the responsibilities, some of which I hadn't anticipated from my experience with smaller meetings—the Salt Lake City, Utah, hotel required me to "guarantee" the room block by signing a contract, for example, which I only later realized would have required me to pay up to $100,000 out of my personal funds if the promised number of participants failed to materialize. The weeks before each meeting were a chaotic series of crises. Over time, though, we set up structures—most notably, establishing program committees for the main meetings and workshops, and hiring a part-time conference coordinator.

Today, the conference runs with only minimal oversight from me and the others on the executive committee. And with input from the program committees and the broader community, we have increasingly had the opportunity to spend more time on self-reflection, considering important issues such as representation and accessibility.

Like many conferences in the early 2000s, COSYNE was exclusively chaired by men for its first five years. In 2008, we shifted to a system in which a man and a woman (one a theorist and the other an experimentalist) would co-chair each year, a first step in our ongoing commitment to diversity and inclusion. COSYNE has also been fortunate to attract sponsorship from a number of nonprofit organizations and corporations, with that funding used primarily for student travel grants.

As COSYNE meets again for the 20th time this week, it continues to evolve. It has swelled to more than 1,000 attendees and outgrown (and been priced out of) ski resorts; The final ski-centric workshop took place in 2023 at Mont Tremblant in Canada. Despite growing interest, we've chosen to limit the meeting's size so that we can continue to foster a single community; the main COSYNE meeting remains a single-track conference that can be housed at one hotel or small conference center.

On the subject-matter side, artificial neural networks, which played a major role in the early history of computational neuroscience but then receded, have returned as a major influence. The chasm between experimentalists and theorists has narrowed considerably, perhaps due in some small part to COSYNE itself. The two groups are, if not exactly bilingual with each other's jargon, far more adept at translation: If you walk into a random talk at COSYNE today, chances are pretty high that some of the slides will be full of equations. But the experimentalists aren't fleeing—indeed, it's quite possible that the speaker showing the equations is an experimentalist.

Experimental

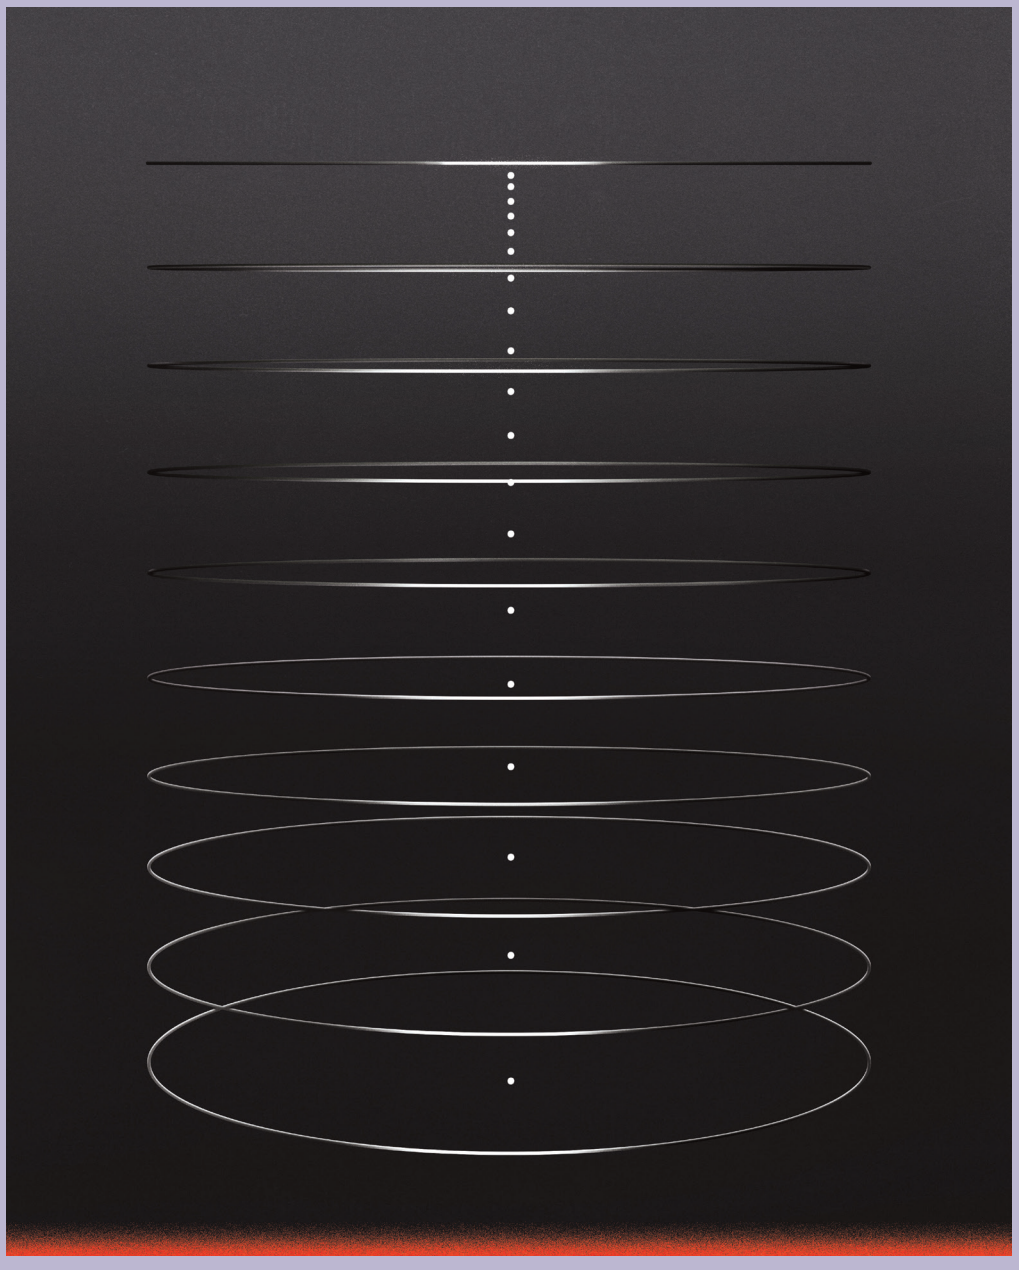

approaches

Learning or performance? Why the distinction matters for memory science

Twisties test: Simone Biles struggled to vault at the Summer 2020 Olympics because she felt unsure of her ability to translate her memory of action into behavior.

PHOTO BY TIM CLAYTON - CORBIS / GETTY IMAGES

New methods make it possible to probe the neural substrates of memory with unprecedented precision. Making the most of them demands careful experimental design.

BY STEPHEN MAREN, DIRECTOR OF THE BECKMAN INSTITUTE FOR ADVANCED SCIENCE AND TECHNOLOGY AND PROFESSOR OF PSYCHOLOGY, UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN

One of the grand challenges in neuroscience is to understand the nature and mechanisms of memory. But memory is elusive. From an experimenter's perspective, memory manifests only in behavior—how we do on a test, for example. But as I write this article, I can recall many memories without any observable change in behavior: My recollection of the amazing sushi dinner I ate last night, for example, exists whether or not I tell anyone about it.

Likewise, losing the ability to do something does not imply that memory is lost. Simone Biles, widely considered one of the greatest gymnasts of all time, did not experience memory loss when she famously got the "twisties" while vaulting at the 2020 Summer Olympics. Her problem was with performance, not memory; she failed to complete the vault because she felt unsure of her ability to translate her memory of action into behavior.

Memory can exist without a behavioral signature. And this concept, known as the learning-performance distinction, has profound implications for how we study the neuroscience of memory. Fundamentally, the learning-performance distinction means that neural manipulations that undermine learned behavior do not necessarily do so by impairing memory. In fact, the behavioral expression of memory can be

influenced by many things, including motivation, fatigue or sensorimotor impairment.

With technological advances that enable scientists to probe the neural substrates of memory with unprecedented precision, this distinction becomes especially important. Yet neuroscientists new to the field are often unaware of the fundamental difference between learning and performance. Failing to appreciate this distinction once central to the neurobiological study of memory has led to work that obfuscates the literature and slows progress.

To advance the science of memory, researchers must design experiments with this distinction in mind. Fortunately, several strategies exist to help researchers tackle the issue.

"The learning-performance distinction, has profound implications for how we study the neuroscience of memory."

Thirty years ago, Richard F. Thompson and his colleagues beautifully dissociated learning and performance using an eyeblink-conditioning procedure in rabbits. In this task, rabbits learn to close their eye when they hear a tone that signals an impending irritating air puff to the cornea. This form of classical (Pavlovian) conditioning is highly adaptive and allows organisms to organize behavior proactively to anticipate future events.

Thompson used an array of tools, including anatomical tract tracing, selective brain lesions and electrophysiological recordings of neural activity, to try to pinpoint the physical substrate of the conditioned response—the memory trace, or engram—in the brainstem. He suspected either the interpositus nucleus (one of the cerebellum's "deep nuclei") or the red nucleus in the ventral midbrain. Lesioning or pharmacologically inactivating either brain region prevented the animals from developing a conditioned eyeblink but did not affect their basic ability to blink (they still responded to the air puff).

The findings suggested either region could be the locus of the elusive engram. To distinguish the regions' roles in learning versus performance, Thompson and his colleagues continued training the rabbits but stopped giving them the inactivating drug. When the interpositus nucleus was awakened, rabbits acquired the task no differently from naive animals that had never received previous training—in other words, the interpositus-inactivated animals seemed not to have learned the task during their previous training.

By contrast, animals initially trained when the red nucleus was inactivated showed perfect retention. They blinked in response to the sound on the very first reconditioning trial, performing identically to animals trained without the drug. Red-nucleus-inactivated animals had acquired the task normally, despite a complete absence of learned behavior during training, showing that the red nucleus is important for performing the learned response but not for learning or remembering it. The memory trace for this behavior,

Thompson concluded, is localized in the cerebellum, a monumental discovery for the field.

Thompson's work highlights the importance of the learning-performance distinction when interpreting the effects of neural circuit manipulations. Had his team not carefully analyzed when memory was preserved in the rabbits, they might have falsely implicated the performance limb of the eyeblink-conditioning circuit in memory processes.

In fact, performance deficits masquerade as memory impairments in many scenarios. For example, memories are best retrieved under the conditions or "context" in which they are initially encoded. Numerous studies in both humans and animals have demonstrated this phenomenon, known as context- or state-dependent memory.

In a textbook example from the 1970s, study participants memorized word lists either on land or underwater (equipped with scuba gear, of course). Minutes later, they tried to recall the lists either in the same context (dry-dry, wet-wet) in which they had learned them or in a different context (dry-wet, wet-dry). Recall was far superior when retrieval occurred in the same context as learning, showing that poor recall of "out-of-context" information resulted from a performance deficit, not from a failure to learn the material.

These effects are not unique to environmental contexts: Drugs (e.g., alcohol or cannabis inebriation), cognitive contexts (e.g., task instructions or strategies), hormones (e.g., gonadal steroid levels) and satiety states (e.g., hunger or thirst) can all yield state-dependent memory, highlighting a confound for many neuroscientific studies of learning and memory—one that is becoming more common as our ability to precisely target brain cells increases. Manipulating neural activity might change the internal context in which information is encoded during learning. When memory is tested later without the brain manipulation, the mismatch in internal context undermines the ability to retrieve the memory.

My laboratory demonstrated just this sort of mismatch. Rats trained to fear a sound when the thalamic nucleus reuniens, which connects the medial prefrontal cortex and hippocampus, was pharmacologically inactivated showed impaired retention the following day. However, we could fully reverse this impairment by inactivating the thalamus during the recall test. Put simply, rats trained and tested in the same brain state retained the learned response perfectly, whereas those trained and tested in different states did not.

Again, our results make it clear that poor performance on a retention test does not imply poor learning or memory loss. The good news is that experimenters can detect performance deficits masquerading as memory loss using methods that reveal when memory is spared. They can identify generalization deficits caused by a mismatch in the learning and retrieval contexts (whether internal or external) by testing animals under common conditions. And learning curves provide important information about the acquisition rate and peak performance of learned responses.

Ultimately, behavioral designs that dissociate learning from performance will be essential to leveraging next-generation technology to advance the science of memory.

On fashion in neuroscience: In defense of freezing behavior

ILLUSTRATION BY NATALIE NELSON

Neuroscience experiments are moving toward the analysis of more complex behaviors, enabled by increasingly sophisticated tools. But we shouldn't abandon simpler paradigms.

BY SHEENA JOSSELYN, SENIOR SCIENTIST, HOSPITAL FOR SICK CHILDREN; CONTRIBUTING EDITOR, *THE TRANSMITTER*

Fashions come and go, even in science. What was once all the rage can be dismissed, viewed as passe over time. In neuroscience, there is a growing appreciation of the importance of studying behavior, enabled by a vast arsenal of tools to manipulate and observe brain function and to track increasingly complex types of behavior. The introduction of markerless pose estimators and machine-learning-based algorithms, for example, enables researchers to automatically quantify even the most complex behavior, from rodent facial expressions to unrestrained naturalistic behavior in the wild, with previously unimaginable resolution.

With this new capacity to study different types of complex behaviors, it may be tempting to view more-established, simple-looking behaviors as archaic. But is it a good idea to blindly follow fashion and relegate such behaviors to the dustheap, akin to last year's skinny jeans? I believe that the brain is best understood by embracing many different approaches, including studying many types of behavior—that there is, and perhaps always will be, a critical place in neuroscience for the study of seemingly simple behavior.

As an example of simple-looking behavior, let's consider freezing, a defensive response to threatening stimuli. Defined as the absence of movement other than breath-

ing, freezing is typically measured in Pavlovian threat experiments as a way to assess learning and memory in rodents. In these experiments, a neutral sensory cue, such as a light or tone, is paired with an aversive foot shock. After learning to associate the previously neutral cue with the shock, the cue alone triggers freezing, known as the conditioned response.

The study of freezing has a long history, dating back to the 1950s and '60s, but came to the fore of neuroscience in the 1990s with brain circuit dissection techniques and the introduction of genetically modified mice. Those experiments revealed, for example, a key brain region in threat conditioning: the lateral nucleus of the amygdala. Researchers mapped how sensory information travels to the lateral amygdala, identified many molecules important in memory formation, including CREB and aCaMKII, and even discovered a new memory phase.

Of late, however, freezing has dropped out of neuroscientific fashion, eclipsed by decision-making tasks and complex naturalistic behaviors. Several explanations may account for this. Freezing was deemed to be a simple, passive behavior that was not ethologically valid—that is, it didn't reflect the animal's natural behavior. Some researchers questioned whether freezing in rodents was a good index of the subjective emotion of fear in people and whether results from these rodent studies could be translated to treat human fear and anxiety disorders. I contend that it is now time that freezing becomes a new retro fashion trend in neuroscience. Here's why:

"Because freezing is a highly motivated behavior, findings gained from conditioned freezing studies have largely stood the test of time and have been reproduced by many labs."

CRITICISM 1: FREEZING IS A SIMPLE, PASSIVE BEHAVIOR.

Despite appearances, freezing is far from a simple, passive response. A freezing rodent is not "lounging." Instead, it shows high muscle tone, reflecting a state of attentive immobility, which may enhance perception and help the animal prepare for a quick escape. Freezing may also help conceal the animal's next move from its predator. Whether or not a rodent freezes is influenced by both internal and external factors, suggesting that freezing reflects the outcome of an active decision-making process. Although not observable by behavior alone, there seems to be a lot going on "under the hood" of a freezing animal.

Assessing freezing behavior in the lab is relatively simple, which I view as a strength. Measuring freezing can be extremely low cost, making this type of experiment accessible to many labs regardless of budget. In the omics era of big data, it might be beneficial to combine complex brain manipulations and observations with a relatively easily quantifiable behavior. Alternatively, researchers can use pose estimators to construct a microstructural timeline of the precise behavioral action sequence of all defensive behaviors, including freezing.

This timeline could be probed by different brain-manipulation and observational techniques.

CRITICISM 2: PAVLOVIAN THREAT CONDITIONING AND THE ANALYSIS OF FREEZING BEHAVIOR ARE ARTIFICIAL "LAB-BASED" PROCEDURES AND BEHAVIORS THAT ARE NOT ETHOLOGICALLY VALID.

Granted, rodents are highly unlikely to encounter electrical shocks in the wild. But rodents do freeze when confronted with threatening stimuli such as a hawk or an owl, indicating that these threat experiments tap into a naturally occurring behavior. Indeed, freezing in response to threats is a natural behavior observed across many species, described in birds, monkeys and people. For instance, freezing behavior (defined as immobility, reduced heart rate and increased muscle tone) has been observed in people watching film clips of car accidents or depictions of social threats, such as angry faces.

Although there is an important place for studying natural self-generated behavior with high ethological validity, examining a motivated behavior such as conditioned freezing in the lab has several advantages. Pavlovian threat experiments afford enormous control, enabling experimenters to minimize the effects of variables that are not of interest and rule out alternative interpretations of their data. Because freezing is a highly motivated behavior, findings gained from conditioned freezing studies have largely stood the test of time and have been reproduced by many labs. And a rich theoretical history supporting conditioned freezing offers neuroscientists the ability to formulate and test hypotheses and constrain findings, thereby increasing the reproducibility of their work.

CRITICISM 3: FREEZING IN RODENTS IS NOT EQUIVALENT TO THE SUBJECTIVE EMOTION OF FEAR IN PEOPLE.

Joseph LeDoux, a neuroscientist at New York University in New York City, has argued that fear is a subjective emotional state in people that cannot be directly measured in (and probably should not be inferred from) other animals, instead suggesting we use terms such as "threat" and "defensive response" in rodent work. I agree. I also agree that, at present, there is a lack of direct translatability of findings from conditioned freezing studies in rodents to human psychiatric conditions. But I argue that this situation is not unique to conditioned threat studies measuring freezing. For instance, though people and rodents obtain food differently, studying foraging behavior in rodents is yielding important insights into brain function. Behaviors studied in rodents need not map 1-to-1 with people to be useful. Our field will progress by gaining a fundamental understanding of how the brain works, and this foundational knowledge, I believe, will be key to more targeted and effective treatments of a myriad of human disorders.

As a now senior neuroscientist, I've witnessed the waxing and waning of fashions in neuroscience. I continue to be excited when new and better methods replace older suboptimal ones. But I believe it is important to continually evaluate the utility of established methods, rather than quickly adopting the new simply because new things are fashionable. It's in this spirit that I am defending freezing behavior; despite its reputation as being unfashionable, it remains an important behavior to gain insights into brain function. To paraphrase Coco Chanel, fashion may change but style endures. Plus, I happen to still look good in last year's skinny jeans.

To make fMRI more clinically useful, we need to really get BOLD

ILLUSTRATION BY MARI FOUZ

A better understanding of the blood oxygen level dependent, or BOLD, signal requires more support for multimodal imaging studies.

Functional MRI (fMRI), though expensive, has many properties of an ideal clinical tool. It's safe and non-invasive. It is widely available in some countries, and increasingly so on a global scale. Its "blood oxygen level dependent," or BOLD, signal is altered in people with almost any neurological condition and is rich enough to contain information specific to each person, offering the potential for a personalized approach to medical care across a wide spectrum of neurological conditions.

But despite enormous interest and investment in fMRI—and its wide use in basic neuroscience research—it still lacks broad clinical utility; it is mainly employed for surgical planning. For fMRI to inform a wider range of clinical decision-making, we need better ways of deciphering what underlying changes in the brain drive changes to the BOLD signal.

If someone with Alzheimer's disease has an increase in functional connectivity (a measure of synchrony between brain regions), for example, does this indicate that synapses are being lost? Or does it suggest that the brain is forming compensatory pathways to help the person avoid further cognitive decline? Or something else entirely? Depending on the answer, one can imagine different courses of treat-

BY EVELYN LAKE, ASSISTANT PROFESSOR OF OF RADIOLOGY AND BIOMEDICAL IMAGING, YALE SCHOOL OF MEDICINE

"We cannot extract sufficient information from fMRI and patient outcomes alone to determine which scenarios are playing out and therefore what we should do when we observe changes in our fMRI readouts."

ment. Put simply, we cannot extract sufficient information from fMRI and patient outcomes alone to determine which scenarios are playing out and therefore what we should do when we observe changes in our fMRI readouts.

To better understand what fMRI actually shows, we need to use complementary methodologies, such as the emerging optical imaging tool of wide-field fluorescence calcium imaging. Combining modalities presents significant technical challenges but offers the potential for deeper insights: observing the BOLD signal alongside other signals that report more directly on what is occurring in brain tissue. Using these more direct measurements instead of fMRI in clinical practice is not an option—they are unethical to use in people or invasive, requiring physical or optical access to the brain. But recent advances in rodent and non-human primate fMRI make it possible to use multiple complementary technologies simultaneously in animal models. Researchers can make tightly controlled observations across an animal's lifespan and test novel intervention strategies, which will be especially important for fMRI to become an effective clinical tool.

Wide-field fluorescence calcium imaging is particularly promising as a complementary technology in that it encompasses a large field of view—in mice, the entire cortical mantle —enabling the researcher to monitor brain activity at the circuit or network-level. It overlaps well with the whole-brain view afforded by fMRI and can be targeted to specific cell types. In combining these two modalities, we can simultaneously observe cell-type-specific activity and the clinically accessible (but indiscriminate) BOLD fMRI signal.

My group has built a device capable of simultaneous wide-field optical imaging and fMRI, which we are using to track the emergence and progression of dynamic changes in functional connectivity in a mouse model of Alzheimer's disease. In research published in Nature Communications in January, we showed that the two methods track well overall, with BOLD signals in line with excitatory neural activity. We also observed some intriguing instances where the two modalities diverged, which may reflect an uncoupling between neural activity and vascular responses.

Our work and that of others hints at the promise of multimodal imaging. But to realize that potential, the field needs to overcome a number of challenges. Obtaining high-quality fMRI data from animals requires an array of skills that go well beyond that required to collect human fMRI data—among them, basic-to-complex animal care (including breeding and genotyping); surgical manipulations and post-op recovery; anesthesia; intubation/extubation and ventilation; animal training (for imaging awake subjects); maintaining animal physiology during imaging; data manipulation and curation; statistics; neurobiology; engineering; disease and injury modeling in animals; machine learning; study coordination; specialized software and hardware development; clinical medicine and inter-species translation; and more. Proficiency in these skills requires years of training and practice, and the number of labs with expertise to give this training is small compared with those doing more traditional human fMRI. For high-quality multimodal data, the demands are even greater, including expertise in both imaging modalities, engineering and their combination.

Adding to the technical challenges are cultural practices that fail to adequately support and credit the people who are performing these difficult and time-consuming experiments, especially when it comes to multimodal animal data. The people doing the work are largely trainees, not technicians, at critical points in their academic careers. The current research environment discourages new trainees from acquiring the necessary skills to collect new data, especially complex data, because it can take years to pay off. Why not just download some that has already been acquired (and work from home)? As a mentor, it is becoming increasingly tough to advise trainees to devote substantial time and energy to data collection.

Sharing data from these types of experiments is also particularly challenging. For fMRI research involving human participants, large, curated and openly shared data amassed from different institutes are an indispensable resource. The same type of dataset for unimodal, and even multimodal, animal imaging would be transformative for fMRI research and the clinical utility of fMRI. This goal requires standardized guidelines for data acquisition, curation and preprocessing, as well as organized and well-managed repositories for data and code. Researchers have begun important work toward establishing these resources for animal imaging data. But some aspects of this undertaking, especially for multimodal animal data, are fundamentally different from the human database counterpart.

Encouraging communication between those who collect data, who often have invaluable information, and the end users will be critical,

"Encouraging communication between those who collect data, who often have invaluable information, and the end users will be critical, particularly at this early juncture."

particularly at this early juncture. This practice will also help to foster relationships among scientists with common interests but diverse backgrounds, which can be a perfect recipe for transformative scientific progress.

To support these types of experiments, the field should invest in and credit the people who acquire these data, grow their ranks, and recognize that generating high-quality data takes skill and time. We also need to increase support for the development of software tools and repositories that enable the proper processing, manipulation, curation and sharing of these data. After all, we have a common goal—a better fundamental understanding of the fMRI signal and better methods for extracting information about the biological processes that underly our measurements. Achieving these ends will have far-reaching impacts on health care by elevating the clinical utility of fMRI.

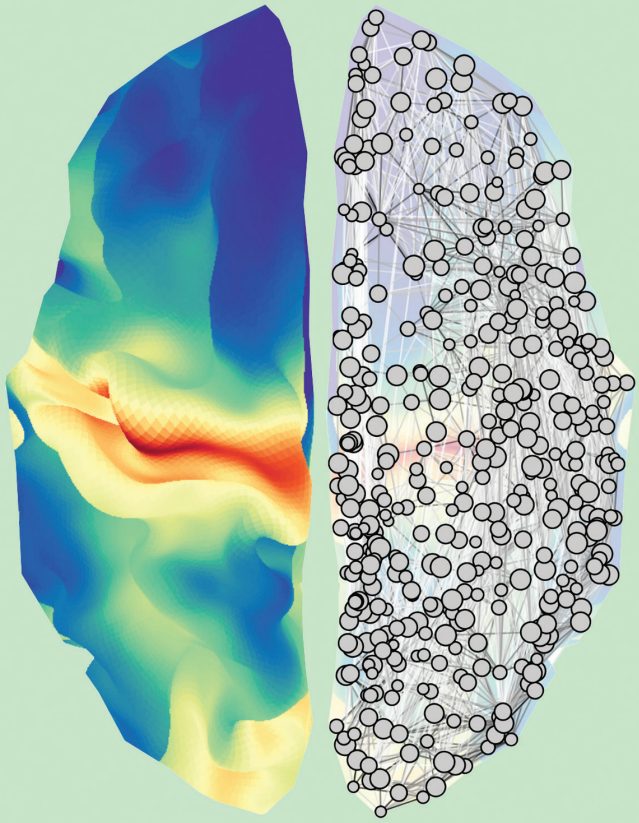

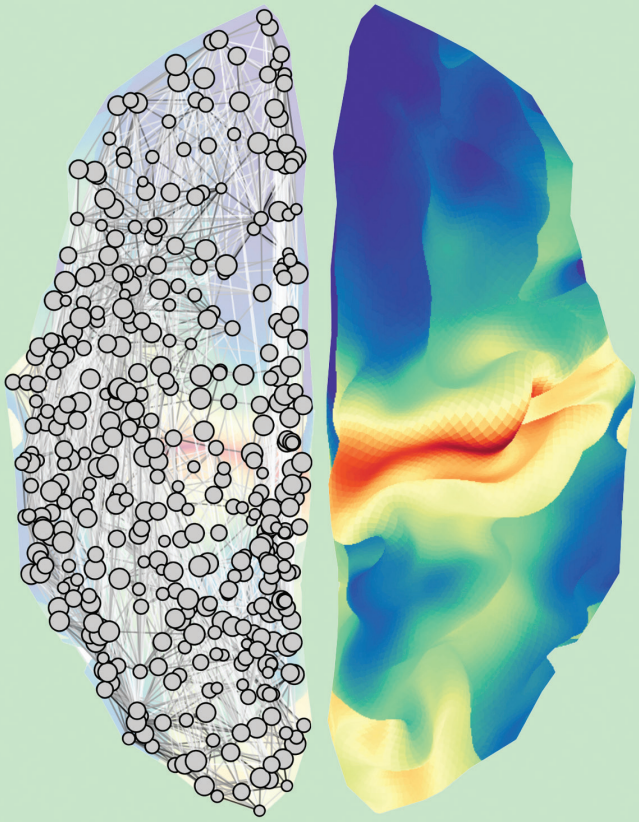

How can we fold cellular-level details into whole-brain neuroimaging networks?

thetransmitter.org/connector-hub/how-can-we-fold-cellular-level-details-into-whole-brain-neuroimaging-networks/

I got answers from Bratislav Misic, who is inventing practical ways to connect the brain's microscopic features with its macroscopic organization.

BY MAC SHINE, ASSOCIATE PROFESSOR OF COMPUTATIONAL SYSTEMS NEUROBIOLOGY, UNIVERSITY OF SYDNEY

BIG PICTURE:

Is the brain uncontrollable, like the weather?

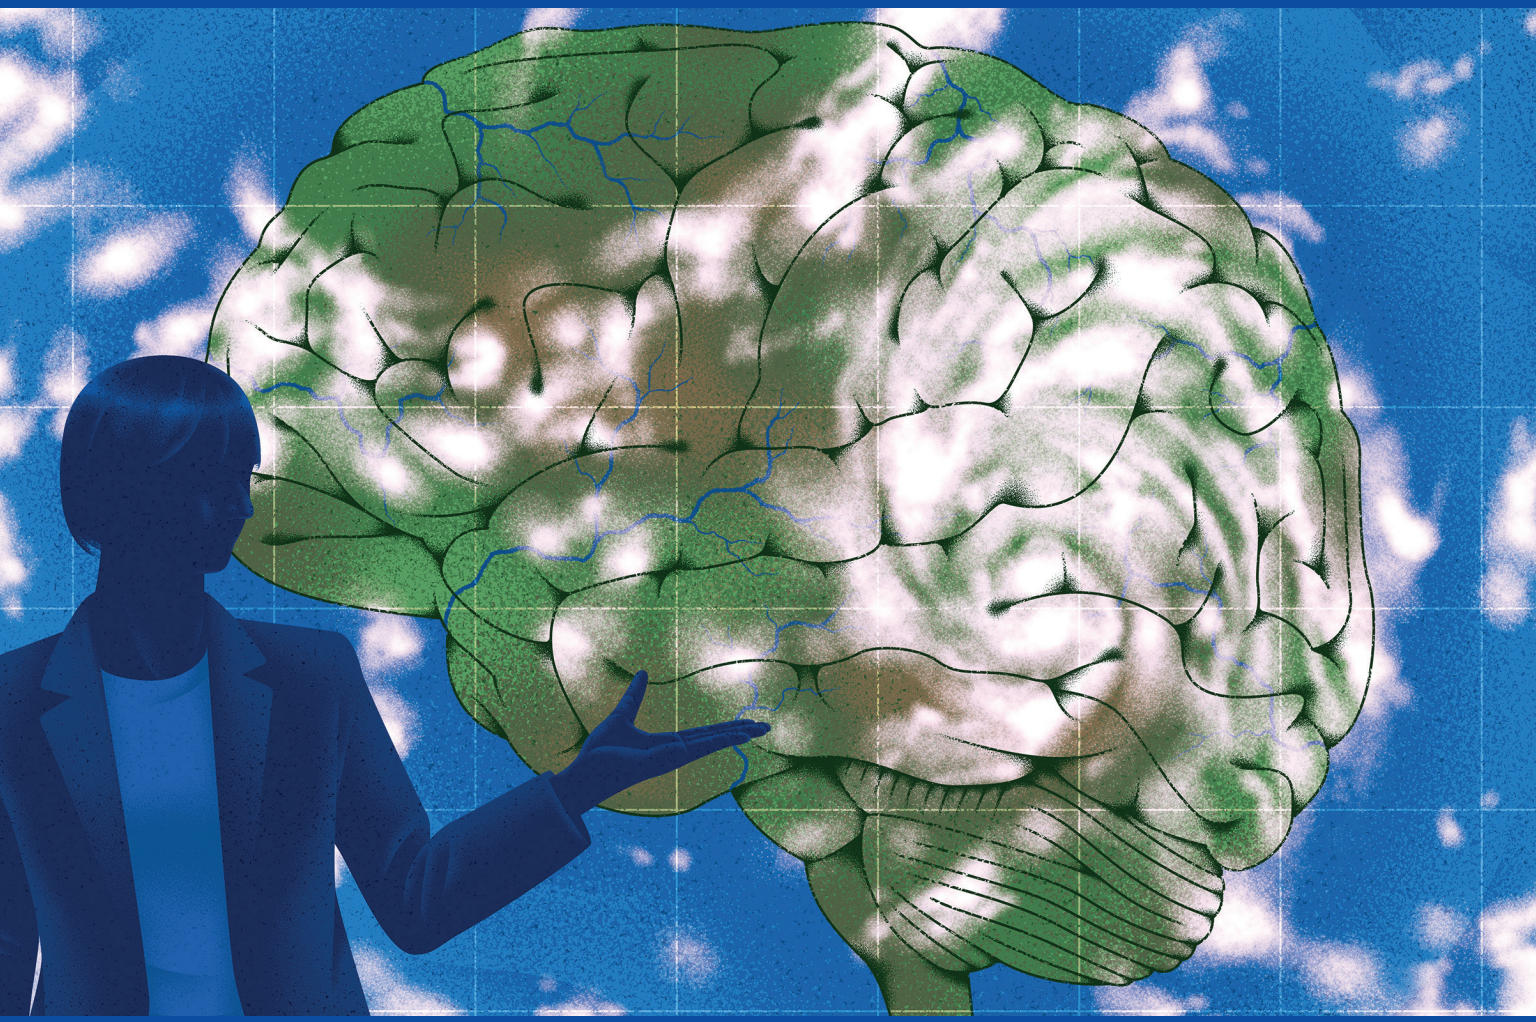

ILLUSTRATION BY KOUZOU SAKAI

The brain may be chaotic. Does that mean our efforts to control it are doomed?

Complex systems operate in ways that are hard to predict from their parts alone, because their behavior is influenced by how their parts interact. As the famous saying goes, more is not just more; "more is different." Brain researchers are increasingly turning to the idea that complex systems support many of the brain's functions, from spatial navigation to memory function. Likewise, they are beginning to realize that many types of brain dysfunction reflect a complex system gone awry, such as when the epileptic brain enters a seizure.

Creating a seizure in a computer simulation is trivial; it's what generally happens when you incorporate excitatory feedback loops with no inhibitory force to counter them. It's vastly more challenging, however, to create a model that approaches the complexity of the brain that isn't seizing. Something akin to this delicate balance is thought to stabilize a number of brain systems, including those that maintain sanity (versus psychosis) and mood stability (versus mania or depression). Indeed, that the brain somehow exists in an exquisite equilibrium the vast majority of the time seems like nothing short of a miracle—given that it relies on numerous giant amplifying feedback loops, offering many avenues to disruption.

BY NICOLE RUST,
PROFESSOR OF PSYCHOLOGY,
UNIVERSITY OF PENNSYLVANIA;
CONTRIBUTING EDITOR,
THE TRANSMITTER

Beyond the brain, many other complex systems live in a similarly delicate balance. Ecosystems can break into toxic blooms. Snow packs can break into avalanches. Weather can break into hurricanes and tornadoes. These systems, however, are not just complex but chaotic, meaning that they are subject to the butterfly effect, by which even tiny perturbations can sometimes push the system out of whack. (The phenomenon gets its name from the way that Edward Lorenz first described it: as if a butterfly flapping its wings over Brazil could cause a tornado in Texas.)

The butterfly effect explains why weather forecasts are much more accurate for the next few days than for the next few weeks. We can measure current conditions with only a certain degree of accuracy, and those small errors in our measurements of what's happening now turn into big errors in our model predictions later.

For the same reason, chaotic systems are exceedingly hard to control, which naturally leads to a question: If the brain is chaotic like the weather—if seizures, depression and psychosis are the analogs of hurricanes—is there any hope of bringing it back to a healthy state via a brain-based intervention, such as a drug or brain stimulation? Or are our efforts to control the brain's complex systems doomed from the outset? In this essay, I contemplate that question. I also asked 14 experts in complex systems to chime in.

Breakthroughs in weather research date back to the early 1900s, when researchers began to formulate the types of weather forecasting models that we use today. Often forgotten is that the explicit goal of early weather research was not just to predict the weather but also control it—both to head off disasters and to weaponize it. Indeed, weather control was the explicit goal behind "the Meteorology Project," organized by Princeton University mathematician John von Neumann and industrial researcher Vladimir Zworykin, the latter of whom contributed to developing the television. As the Second World War began to ramp down in the mid-1940s, the pair approached government officials in Washington, D.C., to request funding for their two-step plan to create a new computing infrastructure to predict the weather (the outcome of which is reflected in today's computers as the von Neumann architecture) and to control the weather using those predictions. As described in their proposal, "Only with exact scientific weather knowledge will effective weather control be possible."

Over the next few decades, other researchers around the globe sought to control the weather. In the United States, a government effort called Project Cirrus, for example, focused on disabling hurricanes. In 1947, the team attempted to dissipate a hurricane, conveniently forecast to remain at sea, by dropping 80 kilograms of dry ice on it from a B-17 bomber. The intent was to disrupt the hurricane's internal structure, but instead the worst possible thing happened: The hurricane's trajectory shifted 130 degrees, and it landed in Georgia. Project Stormfury resuscitated the idea in 1962 and lasted a few decades but never achieved any success. In short, 75 years after the Meteorology Project, we've achieved von Neumann and Zworykin's first goal, forecasting, but weather control hasn't really panned out. Today, weather control happens in subtle ways, such as when China precipitated rain during the 2008

"The brain's complexity opens endless possibilities of creation and reconfiguration of patterns in space and time—and we have no way of predicting when and where they will occur."
—Olaf Sporns

Olympics to ensure that it did not happen at an inopportune time. But because of chaos, we still cannot influence the paths of hurricanes in any predictable way.

For brain researchers, the history of weather control should give us pause. In our own attempts at controlling the brain, how likely is it that we'll face the same difficulties that foiled weather researchers? The answer depends greatly on what type of thing the brain is, and that's still a bit difficult to say. One theory suggests that the brain exists on the knife edge between order and disorder, in a state called "criticality."

This critical brain hypothesis builds on work from physics focused on how phase transitions happen, such as when water changes to steam at a high temperature or carbon changes to diamond at high pressure. In these cases, the large-scale collective property of the system changes when a single parameter, such as temperature or pressure, crosses a critical point. But in other cases, this is not quite the right way to think about it—control derives not from an external parameter such as temperature, but rather from within the system itself. Grains of sand dropped onto a sandpile, for example, elicit avalanches that are just large enough to keep it at the boundary between piling and flattening. Birds in a flock move collectively, but individuals can affect the group's behavior, which is crucial for responding to predators. The system organizes itself in a way that maintains it at that critical boundary of a phase transition.

The critical brain hypothesis accounts for a similar boundary. The gist is that if the brain is too disordered, it can't do anything very useful, akin to being sedated. If it's too ordered, it also can't do anything, akin to being in a seizure. But at the edge of order and disorder, it's optimally positioned to do all the many things it needs to do. The idea follows from studies of criticality in artificial recurrent neural networks, which perform optimally when positioned at the critical boundary. In such networks, the strength of

"The chaos in our brain is a feature we can control and not a maladaptive 'bug' we need to quell."
—Kanaka Rajan

recurrent interactions between model neurons controls where on the spectrum the network sits. If neurons are too interconnected, a small input will trigger every neuron to fire; if neurons aren't connected enough, even a giant input will peter out before it makes its way through the network. But if neurons are connected by just the right amount, an input can and will be processed in a sensible way by a subset of model neurons.

Maintaining the brain at the critical boundary between order and disorder requires some type of exquisite regulation. In artificial neural networks, the balance of excitatory and inhibitory connections maintains criticality—the same may be true in the brain. Along time scales of hours and days, plasticity and other forms of homeostatic regulation, which refers to neurons' capacity to regulate their own excitability relative to total network activity, could help maintain this balance. Along time scales of seconds to minutes, firing-rate adaptation could play this role. Conversely, anything that upsets these mechanisms, including mutated ion channels, broken plasticity mechanisms or aberrant neurotransmission, could throw the brain into a perpetually or partially disordered state.

Although it is a compelling idea, it has been very difficult to test hypotheses of brain criticality. Ideally, we would do things like study how the brain evolves after it is reset to similar initial conditions, and that's just not possible. Instead, most attempts seek to identify the types of signatures typical of systems in a critical state. Phenomena akin to avalanches have been observed in neurons' spiking patterns, for example: bursts of activity in cell cultures that occur with a power law distribution, where small bursts are much more likely than large ones. Another measure relies on the reverberation expected to be triggered in a system, which creates long-range correlations across time. To date, the evidence supports the critical brain hypothesis, but it's far from definitive. We just don't yet know.

Should the brain prove to be chaotic—or close to the critical boundary—what are the implications? Does it mean all hope of control, and therefore treatment, is lost, as is the case for the weather? Or is that the wrong way to think about it? We might consider a few possibilities.

First, some haven't given up on the idea that chaotic systems can in fact be controlled with targeted perturbations. Researchers have figured out ways to control chaotic systems, in theory, via approaches such as the continuous injection of a signal, based on model predictions, as well as perturbations among attractor states. (Attractor states are patterns of activity that a network relaxes into, a bit like a ball rolling into a valley.) But for such an approach to be helpful, researchers would first have to create very extremely precise models of the brain. They would also have to develop ways to control the human brain with much more precision than is typically available today. Under the assumption that this approach would take the form of manipulating either genetic expression or brain activity, it would likely require the control of many genes or stimulation sites.

Second, insofar as disordered states such as seizures are signs of the brain entering subcritical or supercritical states, the brain appears to have internal mechanisms for restoring normal function. Under severe conditions, seizures can continue for hours—illustrating that the brain is physically capable of it—but typically last just minutes. Likewise, people often enter depressive and psychotic episodes and then exit them days or weeks later. Unfortunately, we don't really understand the mechanisms by which the brain self-organizes and renormalizes. A better understanding of those mechanisms could lead to better treatments or preventions, akin to the fences used to prevent avalanches.

Of course, the answer may be the one we wish were not true: It may be that in some cases, we simply cannot control the brain—at least not in the ways we would need to treat some types of dysfunction, such as epilepsy, psychosis and depression.

What do researchers predict?

To get a sense of how likely it is that the brain will turn out to be uncontrollable, like the weather, I asked some experts in complex systems to chime in. Read their responses beginning on the following page.

What do researchers predict?

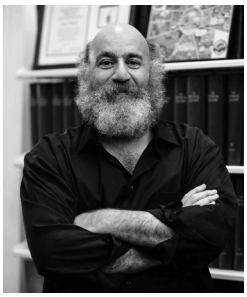

WILLIAM BIALEK, PRINCETON UNIVERSITY

It seems plausible that important functions of the brain emerge from interactions among many neurons. The "more is different" idea then leads to many ways in which neural activity could be unpredictable or uncontrollable. In an ordinary magnet, we can push on the important collective modes of the system just by applying a magnetic field. But even in older, relatively simple models of neural networks, the analog of a magnetic field would be a complicated combination of inputs to each cell in the network, exciting some of the cells and inhibiting others. We don't really have experimental tools that allow us to do this. So, even before we get to more controversial ideas such as criticality or chaos, we have serious problems.

MANLIO DE DOMENICO, UNIVERSITY OF PADUA

Controlling a complex system such as the human brain is a formidable and challenging task. Evolutionary forces have done an extraordinary job of shaping the structure and dynamics of the brain: It self-organizes in response to internal and external perturbations through mechanisms that are still not fully understood. From a statistical physics perspective, unraveling how the brain—as well as other biological systems—is able to self-regulate its behavior while self-correcting for localized dysfunctions might open the door to a plethora of applications for systems biology and systems medicine in general, a perspective that makes the future rather exciting.

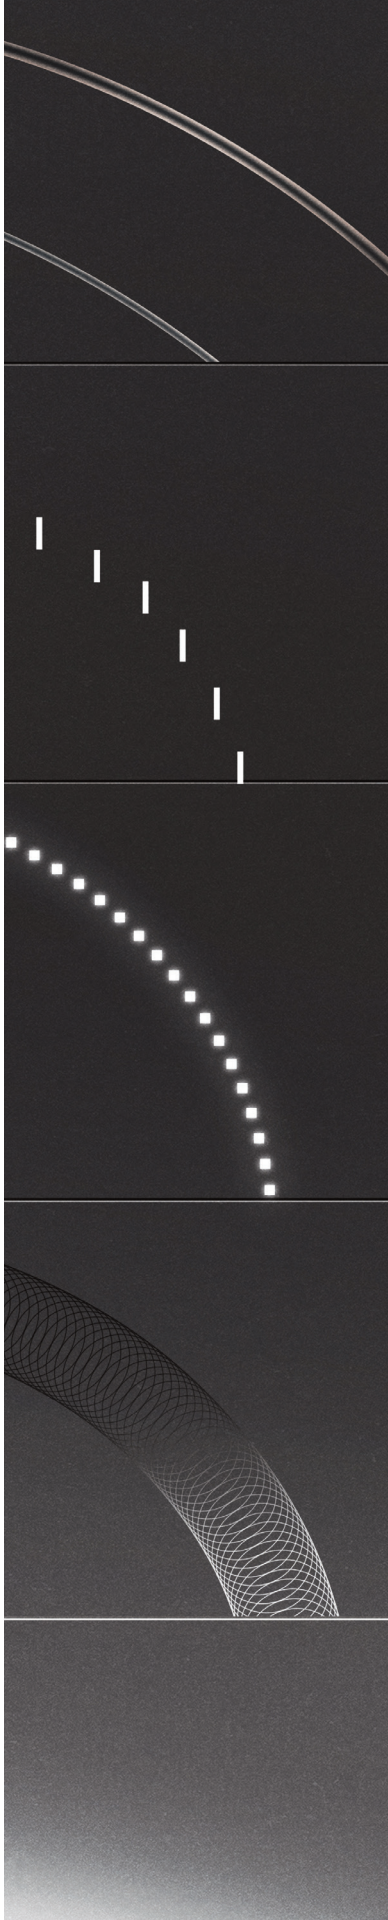

SHAUL DRUCKMANN, STANFORD UNIVERSITY

My intuition, and it is just an intuition, is that the brain will be controllable (by controllable I mean something like pushing the brain out of an epileptic state). The main reason for that is that it needs to be internally controllable—information from one area needs to be transferred to another, which can be thought of as one brain area controlling the state of another. Such forms of internal control and modulation are a core part of how I imagine brains work. If we could tap into something similar to these internal dials, we have a chance at controlling the brain. Our modes of access will be quite different, however. How to use that access will require a sophisticated understanding of how to influence complex systems, which we currently just don't have. Moreover, more subtle control, such as correcting just the parts of the dynamics that changed as the result of neurodegeneration, for instance, similar to diverting but not scattering a hurricane, would require a deeper understanding still. The problem itself is not foreign to science—controlling dynamical systems is a rich field of engineering—but it tends to focus on more straightforward engineered systems, not the complex web of interactions among heterogenous units that is something like a brain. This is exactly the kind of challenge that I and many others are going to devote a few decades of our lives to understanding.

TATIANA ENGEL, PRINCETON UNIVERSITY

The brain differs from the weather in a way that may make it more amenable to control. The brain contains endless self-organizing, self-regulating loops, acting ceaselessly to tune it to a well-functioning state. This self-tuning property is common to all complex biological systems. Think about how your body maintains a nearly constant temperature over a broad range of ambient conditions. Similarly, the brain has mechanisms across all scales—from molecules to large-scale networks—for sensing deviations from the normal and steering itself back to the functional state. Many disease states result from a malfunction in one of the self-tuning mechanisms. Thus, if we could fix the self-tuning mechanism, restoring the production of a missing molecule, for example, this repaired mechanism would almost miraculously do all the hard work of steering the brain toward the functional state. Although this general

idea sounds simple, it may be hard to realize in practice because the many self-tuning mechanisms are highly intertwined. The same molecule may be part of several self-tuning loops, and in trying to repair one, we could destroy another. In addition, the internal compensatory mechanisms may make a diseased brain different from a healthy one. For example, when one brain area is lesioned, another area can take over its function. Or an area deprived of inputs necessary to perform its function can take on a new function, such as when a visual area in a blind person responds to sound or touch. Given these kinds of alterations, the same self-tuning loop may produce different outcomes in a healthy versus a diseased brain. Technological advances enable us to interface with the brain with increasing spatial and temporal precision, but the problem of controlling the brain remains far from being solved.

STEPHANIE R. JONES, BROWN UNIVERSITY

I don't think the brain is too complex to control to aid in the treatment of neuropathology. There are many examples showing that noninvasive electrical and magnetic perturbations to the human brain can help restore normal function. One area in which there has been particular advance is in the treatment of depression, where regular patterns of stimulation are improving symptoms by "renormalizing" circuit function. Single pulses of brain stimulation combined with electrophysiological recording of the brain's response are also used to measure the complexity of the response in people in a coma, as a means to predict recovery. But our understanding of how these perturbations directly affect human brain circuits, and if and when they will have lasting effects, is limited. One approach to improving this understanding is by building detailed dynamical models of the biophysical elements that generate electromagnetic activity in cells and circuits, and simulating their response to various patterns of stimulation. With this approach, we can better understand the nature of the brain's complexity and ultimately use it to our advantage to designe more efficacious stimulation paradigms.

ANN KENNEDY, NORTHWESTERN UNIVERSITY

I wouldn't necessarily despair about unpredictability. The Lorenz system is a famous example of a chaotic attractor, developed as a model for atmospheric convection (to keep with your climate theme.) You can either change the 3D state (x, y, z) of the system or you can change the three quenched parameters (sigma, rho, beta) that govern how it evolves over time. The latter is indirect but much more powerful, even allowing you to eliminate chaos in the system entirely. Though I can't tell you where the state of a Lorenz system will be far into the future, I can tell you that if you keep its quenched parameters fixed, its state will always lie within a tiny fraction of the total volume of 3D space. I believe this metaphor will hold for the brain: If we can create targeted interventions that speed, slow, amplify or suppress the flow of neural activity through particular brain regions, we can restrict the space of trajectories activity will take through those regions even without controlling said activity directly. This is what I think a lot of neuromodulators and neuropeptides are doing: reshaping the neural substrate to direct the flow of fast patterns of neural excitation and inhibition. Our challenge is that the brain has had millennia of evolution to make sure the right reshaping signals go to the right bits of substrate—we need to understand what it is they are doing there so we can hope to create interventions that mimic them with the same selectivity and specificity as the brain itself.

EVE MARDER, BRANDEIS UNIVERSITY

I think of these problems quite differently. Instead of focusing on the fact that diseased brains or normal brains that are faced with extreme perturbations "crash," I would like to emphasize that the brain has many cellular and molecular mechanisms that promote stability. Just because it is possible to trigger brain dysfunction shouldn't lead one to think that all brains are teetering on the edge of dysfunction. Rather, there are multiple sets of cell and circuit parameters that are consistent with "good enough" behavior, and this allows circuits to wander around in parameter space without losing function. And there are numerous and overlapping mechanisms that support cellular and neuronal stability. My lab studies the effects of temperature and other extreme perturbations

on the crustacean stomatogastric nervous system. Although all animals will "crash" if you raise the temperature enough, they are resilient to the more than 20 degree Celsius temperature fluctuations that they usually experience. Many mechanisms play a role in this resilience, and likewise, many mechanisms also play roles in the resilience of healthy human brains.

MARINO PAGAN, UNIVERSITY OF EDINBURGH

Unlike a hurricane's dynamics, those of the brain are subject to powerful regulatory mechanisms and are finely honed during development to perform specific functions. These compensatory forces ensure that most brains don't erupt in seizures, and that relatively normal function can be maintained even in the presence of lesions or genetic mutations. I believe that a deeper understanding of the brain's intrinsic regulatory mechanisms will be crucial to learning how to provide corrections when neural dynamics enter unhealthy states. Unfortunately, no two brains are alike, and I suspect that such an understanding will need to be tailored to individual brains, and that "precision medicine" approaches will be necessary to recapitulate the large degree of individual variability that accompanies almost every type of brain dysfunction. One hope, however, is that some of the key mechanisms will be best described in a "latent space," abstracted away from the extraordinary complexity of neural circuits, and that such high-level descriptions will be more amenable to scientific inquiry and to treatments.

STEPHANIE PALMER, UNIVERSITY OF CHICAGO

The brain's initial processing, the sensory "shell" that takes in and processes input, may have the most to gain from being poised near criticality. That gives the system exquisite responsiveness to changes in the external world. Whether that's architected by the brain or a consequence of the structure of the driving input is up for debate, but the fact remains that signatures of criticality are observed in many sensing systems. Deeper in the brain, I'd expect neural populations to be, if anything, less critical—further from this kind of singularity. (Though, of course, hippocampal regions are the cradle of epileptic foci

in humans.) What is clearly true is that the brain functions most of the time, in most organisms, in most individuals. This likely means that it's not such a knife's edge. As a theorist, I hope this feature of neural coding winds up being understandable and interpretable, not just a collection of patches that biology implemented over evolution's punctuated trudge through functional space. Of course, biology doesn't owe anything to a theorist. But I hope evolution also found some controllable knobs that we can discover.

LUIZ PESSOA,
UNIVERSITY OF MARYLAND, COLLEGE PARK

My guess is that the brain is not as uncontrollable as the weather! In my view, the brain is a complex system that works in a highly distributed, heterarchical manner (i.e., lacking a clear hierarchy). Because it works in an integrated way with the body and environment, brain signals circulate in ways that are extremely hard to predict. But I wouldn't go as far as viewing it as uncontrollable because of the inextricable link between brain and life. Brain and body systems are self-maintaining and in a continual process of homeostasis, which stabilizes their dynamics to remain within bounds that are compatible with life.

KANAKA RAJAN, HARVARD UNIVERSITY

All cognition is dynamic, and the engine that produces cognition—the brain—is a complex dynamical system. Elegant theories based on the physics of large networks of idealized neurons, some of which I am proud to have written myself, have modeled the brain as a chaotic system. (Although, unlike the Lorenz attractor, which is a low-dimensional system described by three variables, neural activity in brains is thought to be more consistent with high-dimensional chaos.) In the face of all this chaotic activity produced internally by neural circuits in the brain, how do we manage to think or behave cogently? It turns out that neural circuits can actively turn down their intrinsically generated chaos when paying attention to even subtle sensory inputs. Interestingly, this mechanism—or phase transition—was theorized first and then verified experimentally through recordings from a number of brain areas. I think that this ability sup-

press intrinsic chaos is due to how we can think and behave, avoiding both hallucinatory oblivion and reflexive entrainment to our inputs or environment. This is just one of the ways by which the chaos in our brain is a feature we can control and not a maladaptive "bug" we need to quell.

CHRISTOPHER ROZELL, GEORGIA INSTITUTE OF TECHNOLOGY

The brain is an enormously complex system, and (like the weather) it may be impossible to exert precise control over it at scale. But is that level of control necessary? Extending the weather analogy a little further, we build houses with a locally controlled environment that keeps the temperature in a comfortable range. Clinically, we have multiple examples of neuromodulation approaches where targeted stimulation doesn't "control" the brain precisely but still influences brain state enough to reduce tremor, seizure activity or depressive moods. Scientifically, approaches such as optogenetics have produced meaningful insight even in their basic form. Empirically, it seems we can make meaningful clinical and scientific progress without full control of the brain.

OLAF SPORNS, INDIANA UNIVERSITY

Is the brain the sort of system that is predictable or controllable? My answer would be: It depends on what you mean by "prediction" and "control." There are certainly many examples of a specific perturbation having a predictable effect or outcome, and there has been strong emphasis in neuroscience on characterizing such cause-effect couplings. But the brain is more than that, a growing realization that requires a shift in thinking and perspective. I tend to approach the brain as a complex system consisting of a huge number of interconnected elements or neurons. When those elements become active, their individual states become entangled or mutually dependent, thus creating high-dimensional informational structures that we are just beginning to glimpse and understand. This collective action of many of the brain's elements results in a continuous flow of activity that underpins cognition and behavior. The system has some level of predictability—consistent topography, patterns of synchrony

and dimension reduction, for example. Brains are far from random, and this allows recognizable features to emerge. But, as for all truly complex systems, there's a limit to prediction and control. For example, predicting specific brain states far in advance, similar to forecasting the weather over a period of weeks or months, is fundamentally impossible, as nonlinearities and chaotic fluctuations quickly take over. For those who want absolute control, this may seem an insurmountable challenge. But I see in it something far more comforting, even liberating. The brain's complexity opens endless possibilities of creation and reconfiguration of patterns in space and time—and we have no way of predicting when and where they will occur.

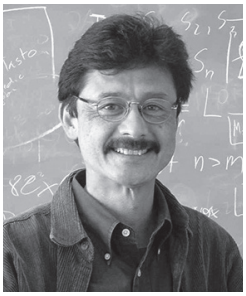

GEORGE SUGIHARA,
UNIVERSITY OF CALIFORNIA, SAN DIEGO

As an ecologist and a wannabe neuroscientist, I think this is a super exciting area for dynamical systems thinking. Clearly the genie is out of the bottle with respect to prediction and understanding, and perhaps with regard to control as well: Witness recent results that try to figure out which of many possible brain-stimulating electrodes will cause the right effects in ways that circumvent the need to try them one by one. Here, researchers predicted the effects of targeted brain-area stimulation from resting activity using a nonlinear dynamical causality test (convergent cross-mapping). It's a great step, and beyond it there's so much untapped potential.

For instance, being chaotic (and thus nonlinear) means that you can't really think of the individual parts of a dynamic system as being separate. This is what mathematicians would call "nonseparable," meaning you can't formally study one piece independently of others. Flipped around, this interdependence has huge advantages because with nonlinear dynamics and chaos in general, any one part of a dynamic system can have information about all of the other parts. This enables us to recreate a shadow version of the whole system from just one piece if it. In ecology, this has enabled us to predict future states of systems such as salmon populations, even when we don't have access to measures of all the causal variables. Taking this a step further, because multiple shadow versions are possible, the same information can be represented simultaneously in factorially different ways, which is something one might imagine in how the brain works. I anticipate that this and other dynamical systems approaches will figure prominently in the quest to understand and treat the brain.

Data-sharing and

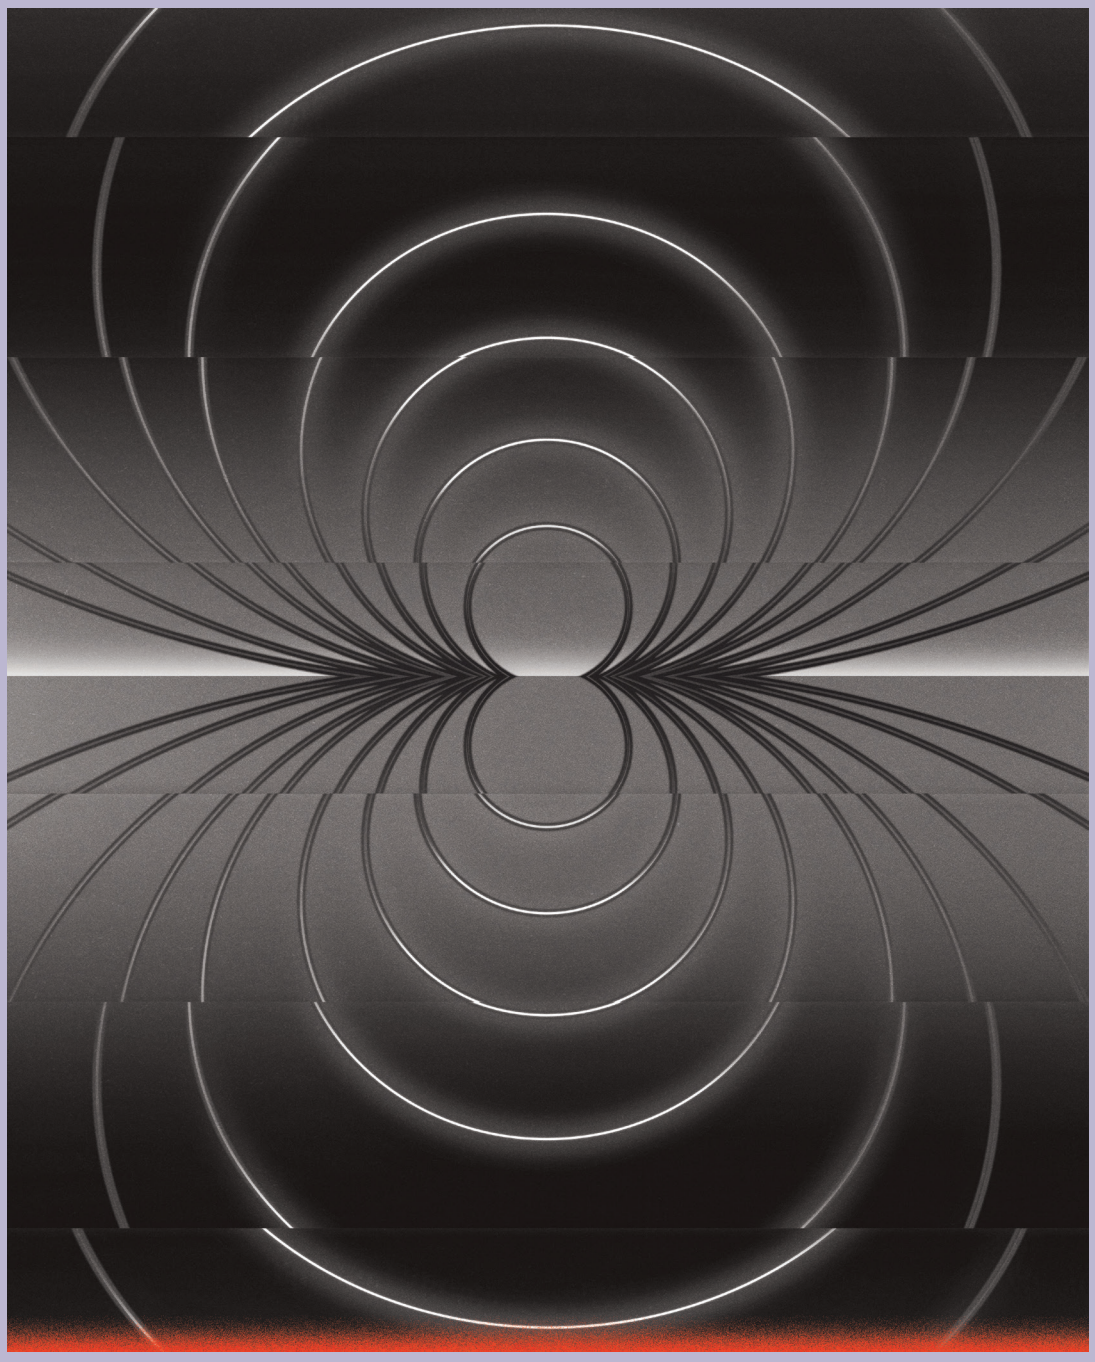

open neuroscience

Neuroscience is in the midst of a major culture shift. Data, once a private asset to be mined only by its creators, is rapidly becoming a resource to be shared. To increase data accessibility and reuse, U.S. federal funders have required since January 2023 that grant applicants include a plan for how they will manage and store their data. Though widely welcomed, the move has come with growing pains. This series of scientist-written essays explores some of the benefits and challenges of data-sharing that researchers have encountered along the way.

Incentivizing data-sharing in neuroscience: How about a little customer service?

ILLUSTRATION BY DANIEL LIÉVANO

To make data truly reusable, we need to invest in data curators, who help people enter the information into repositories.

A growing number of funder and journal policies now require neuroscientists to share the data they produce. But to truly reap the benefits, data must be shared well—no easy feat in neuroscience, which generates a multiplicity of complex data types across spatial scales. What will motivate neuroscientists to spend the time and resources required to make their shared data truly useful to others? Because without adequate support, the new data-sharing policies are in danger of becoming simply a "box-checking" exercise.

If we really want to change the culture around data management and sharing, we must focus on improving the data submission experience. In these early, crucial days of wide-scale data-sharing, we must encourage repositories to take a more customer-service-oriented approach to data submission and offer the necessary support to do so. To be most effective, repositories must work to guide and assist researchers through the submission process and not merely point them to documentation or provide feedback about where they failed.

Thanks to investments in large brain projects and organizations such as the International Neuroinformatics Facility, neuroscience has many of the pieces in place to make shared

BY MARYANN MARTONE, PROFESSOR EMERITA, UNIVERSITY OF CALIFORNIA, SAN DIEGO

"We need to deploy professionals, such as knowledge engineers and data curators."

data useful. Researchers have developed a set of guidelines for data reuse, known as "FAIR": findable, accessible, interoperable and reusable. And neuroscience-specific repositories exist to serve specific data types, neuroscience domains or geographical regions. In recent years, these investments have also spurred the development and adoption of neuroscience-specific standards such as the Brain Imaging Data Structure (BIDS), Neurodata Without Borders, the National Institutes of Health's Common Data Elements and Common Coordinate frameworks. By supporting these standards, repositories are seeding an ecosystem of tools around particular data types.

But for this infrastructure to pay off, data producers have to be both willing and able to populate these resources with standardized, well-curated data. The fact remains that preparing data for a repository, especially one that requires strict adherence to data and metadata standards, is a significant burden that falls asymmetrically on the investigator. If the burden is perceived to be too great relative to the rewards, researchers have the option to go elsewhere, such as a generalist repository with few submission requirements. Indeed, some repositories have lowered their requirements to ensure they still have customers, but the data they house are far less useful.

Fully preparing data for publication requires skills that researchers rarely possess. Data producers typically have a deep understanding of the data but not the mindset, knowledge or resources to adhere to data and metadata standards that optimize data for reuse. Researchers often submit poor-quality metadata, for example, as work from my group and others has shown. I don't believe that simply providing researchers with better training in data management or data science will solve the problem.

Instead, we need to deploy professionals, such as knowledge engineers and data curators. Curators can format and document data according to the standards in place at a specific data repository. They can review submitted data, engage with the submitters where necessary to ensure compliance and often provide additional services, such as mapping metadata to controlled vocabularies or tagging data with keywords.

A few neuroscience repositories, such as EBRAINS (previously the European Human Brain Project Neuroinformatics Platform), Stimulating Peripheral Activity to Relieve Conditions (SPARC), the Open Data Commons for Spinal Cord Injury and the Open Data Commons for Traumatic Brain Injury, have blazed a trail and already invested in curators to improve the consistency and quality of submitted data. Informal surveys show that investigators

are often surprised by the positive impact curation has on their work—their data are now "FAIR," not only to the community, but to the originating lab. When a postdoctoral researcher leaves, their data can be reliably found, accessed and understood. And as researchers work with curators, they start to appreciate how the requirements of a repository—including the use of identifiers, metadata, specific standards and data dictionaries—serve data management overall, and they begin to develop practices within their own laboratories to facilitate sharing.

If researchers are supported properly, submitting their data to specialized neuroscience repositories provides practical training to empower effective sharing across the data lifecycle. It catalyzes a feedback cycle in which benefits, tools and knowledge flow back from the repository to the submitting laboratory and back out to other users, whether human or artificial intelligence, in the form of higher-quality, FAIR data. Everybody wins.

My experience suggests this approach is effective. SPARC has some fairly stringent data requirements; because the NIH project collects high-quality and varied data on the interaction of the autonomic nervous system with end organs, it uses a cross-modality data standard called the SPARC Data Structure, based on BIDS, to organize the variety of data submitted. In the early days, scientists submitting data found the process frustrating and labor intensive, leading to many angry emails. While the technical team worked to improve the infrastructure, the curators worked to establish good relationships with the investigators, acknowledging when the process was difficult and assisting them over any barriers. Over time, curators observed that establishing a respectful, supportive relationship with data submitters rendered their experience much less burdensome. And despite early significant frustrations, when surveyed, many investigators indicated that they intend to continue using SPARC as their data-sharing platform even after their SPARC-specific funding ends.

Human curation is expensive and hard to scale, and funders are often reluctant to pay for it. But I don't believe that this level of human support will be needed forever. In the SPARC project, for example, a young investigator, Bhavesh Patel, seeing the effort required to organize and upload data, developed a software wizard called SODA to automate file-level operations and to guide the researcher step by step through the process. As researchers started to understand what was being asked of them and began using SODA, the process became more efficient for both submitters and curators. We can expect the rapid advance of AI to have a significant impact on curation, data integration and other data challenges.

But in the meantime, we need good data, and that will come from well-curated, standardized and well-managed data in specialist repositories. Reducing the burden on the data submitter, not by lowering requirements but by investing in customer-service-oriented curation, will go a long way toward unleashing the full power of data science on this most complex of organs, the brain.

Disclosure: Martone is on the board of directors and has equity interest in SciCrunch Inc, a tech startup out of the University of California, San Diego that develops tools and services for reproducible science.

Simply making data publicly available isn't enough. We need to make it easy—that requires community buy-in.

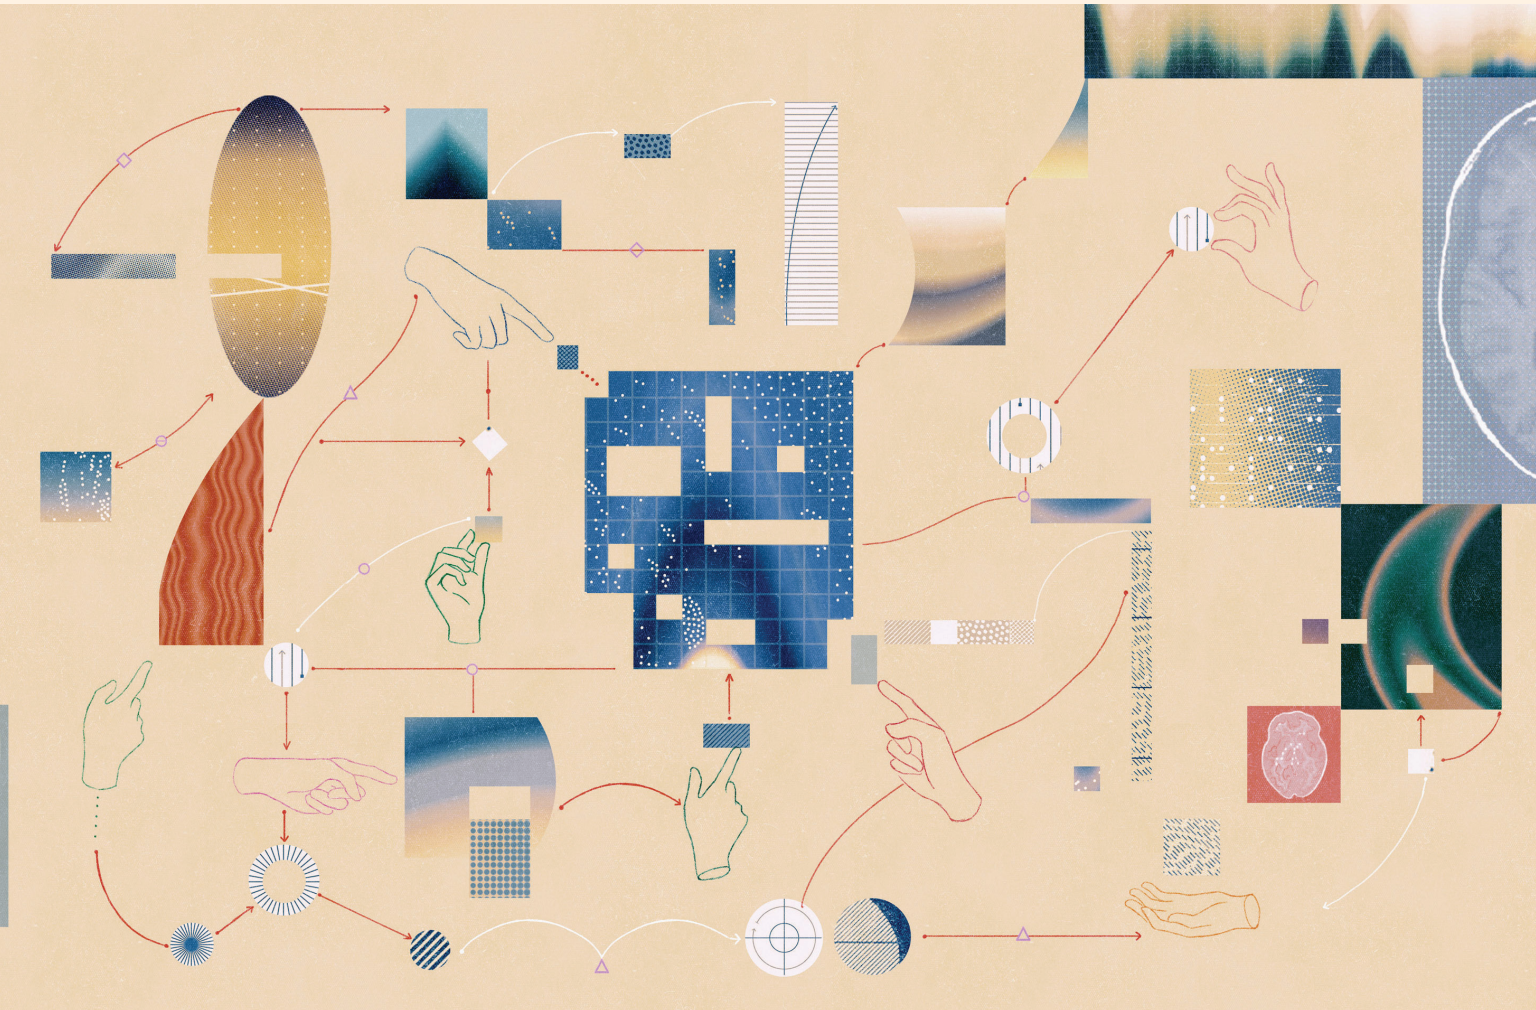

ILLUSTRATION BY DANIEL LIÉVANO

I helped create a standard to make it easy to upload, analyze and compare functional MRI data. An ecosystem of tools has since grown up around it, boosting reproducibility and speeding up research.

The sharing of data collected in scientific studies is increasingly viewed as an important way to make science better. It helps maximize the benefits that accrue from the data (which are often collected using taxpayer funds); it enables larger studies by aggregating across many smaller datasets; and it enables researchers to check the results from published papers to see if they hold up, using the same or different analysis methods.

But shared data provide such benefits only if they are organized in such a way that other researchers can use them effectively. A humorous YouTube video, "Data Sharing and Management Snafu in 3 Short Acts," by the NYU Health Sciences Library in New York City, demonstrates how data-sharing can go wrong. In this animated video, one researcher (played by a sad panda) requests data from another researcher from their paper recently published in the journal Science. After much back and forth, the researcher finally provides the data, leading to the following exchange:

— *I received the data, but when I opened it up it was in hexadecimal [an indecipherable digital format].*

— *Yes, that is right.*

BY RUSSELL POLDRACK,
COGNITIVE NEUROSCIENTIST,
STANFORD UNIVERSITY

— *I cannot read hexadecimal.*

— *You asked for my data, and I gave it to you. I have done what you asked.*

My group has experienced firsthand the importance of data organization. About a decade ago, inspired by Michael Milham and his colleagues in the 1000 Functional Connectomes Project, we started sharing brain imaging data through a project called OpenfMRI. In that early project, researchers sent us datasets to be shared, and a data curator in our group had to reorganize their data to match our in-house organization scheme. This reorganization often required extensive back-and-forth between the curator and the owner of the data.

In 2015, we received funding from the Laura and John Arnold Foundation to expand this project, which ultimately became the OpenNeuro data archive, an open platform for sharing magnetic resonance imaging (MRI), positron emission tomography (PET), magnetoencephalography (MEG) and electroencephalogram (EEG) data. Because we wanted to be able to accept data broadly without requiring a large team of curators, we decided to develop a data organization standard that any researcher could use, enabling them to upload their data without the need for human curation. We realized that this would be successful only if we got many researchers in our community on board, so we worked with a large number of people to develop a new framework that we called the Brain Imaging Data Structure (BIDS). It took about a year to develop the first version of BIDS, which was published in the journal Scientific Data in 2016.

> "By various estimates, there are hundreds of thousands of datasets in the wild that have been converted into the BIDS format."

So what is BIDS, exactly? It's really two things. First, it's a scheme for naming and organizing the many files that make up a brain imaging dataset. My lab uses MRI. BIDS tells us how to name the files that are generated during brain imaging experiments and how to set up the folders that the different kinds of data will go into. BIDS also provides a scheme for how to organize the metadata that describe how the data were collected. Each image file in a BIDS dataset has an associated file that contains detailed information about how the image was collected.

Importantly, BIDS specifies the vocabulary that can be used to name each parameter. For example, "repetition time" is an important variable in MRI experiments, and in the literature it is referred to in many different ways—"Repetition Time," "RT" and "TR," for example—and can be expressed in seconds or milliseconds. BIDS dictates the specific term used to define this value in the metadata ("RepetitionTime"), as well as the units (seconds). To anyone who isn't an MRI aficionado, this level of detail probably sounds immensely boring, but it turns out to be essential if humans or machines are to read MRI datasets without any ambiguity about how the data were generated.

Part of BIDS' success stems from its strong community-driven character: Decision-making is led by an elected steering group, and its ongoing development and maintenance is supported by a group of nine volunteers, along with many other contributors. This kind of community organization requires a lot of time and effort and a willingness to compromise for the greater good.

The BIDS community has grown to be very large, with several hundred researchers having contributed to the effort in some way. By various estimates, there are hundreds of thousands of datasets in the wild that have been converted into the BIDS format, with data from more than 30,000 people available via the OpenNeuro archive alone. The ease of reusing a BIDS dataset has led to many published reuses, helping to maximize the benefits of those data for the community and the world.

In one notable study designed to assess reproducibility, 70 groups of researchers analyzed the same large functional MRI (fMRI) dataset distributed in the BIDS format. The results varied widely depending on analysis workflows, highlighting the need to understand how analytic variability affects scientific results. BIDS made it possible for each of the participating groups to take the dataset and immediately understand how to process it; without BIDS, the degree of communication required to explain the data would have been overwhelming for such a large number of groups.

Yet another great benefit of BIDS is the ecosystem of tools that has grown around it. This suite of community-generated "BIDS Apps" makes it easy to process the data in various ways. These apps enable users to take commonly used imaging analysis software, such as the FreeSurfer tool for anatomical processing, and easily apply it to their BIDS dataset, rather than having to reformat the data to meet the distinct requirements of each software package. One such BIDS App for the preprocessing of fMRI data, fMRIPrep, has become remarkably popular, with several thousand uses each week over the past year. Because BIDS Apps are packaged with all of the required additional software libraries, they also provide a greater degree of reproducibility across different computer platforms.

BIDS also has a defined process for extensions into new data types, which supports growth into new communities. Community members can propose extensions that can then be developed through discussion with the maintainers and the steering group. This process has led to support for many new data types, ensuring the continued growth and relevance of the standard, and demonstrates the strength of the community model for data standards.

How scuba diving helped me embrace open science

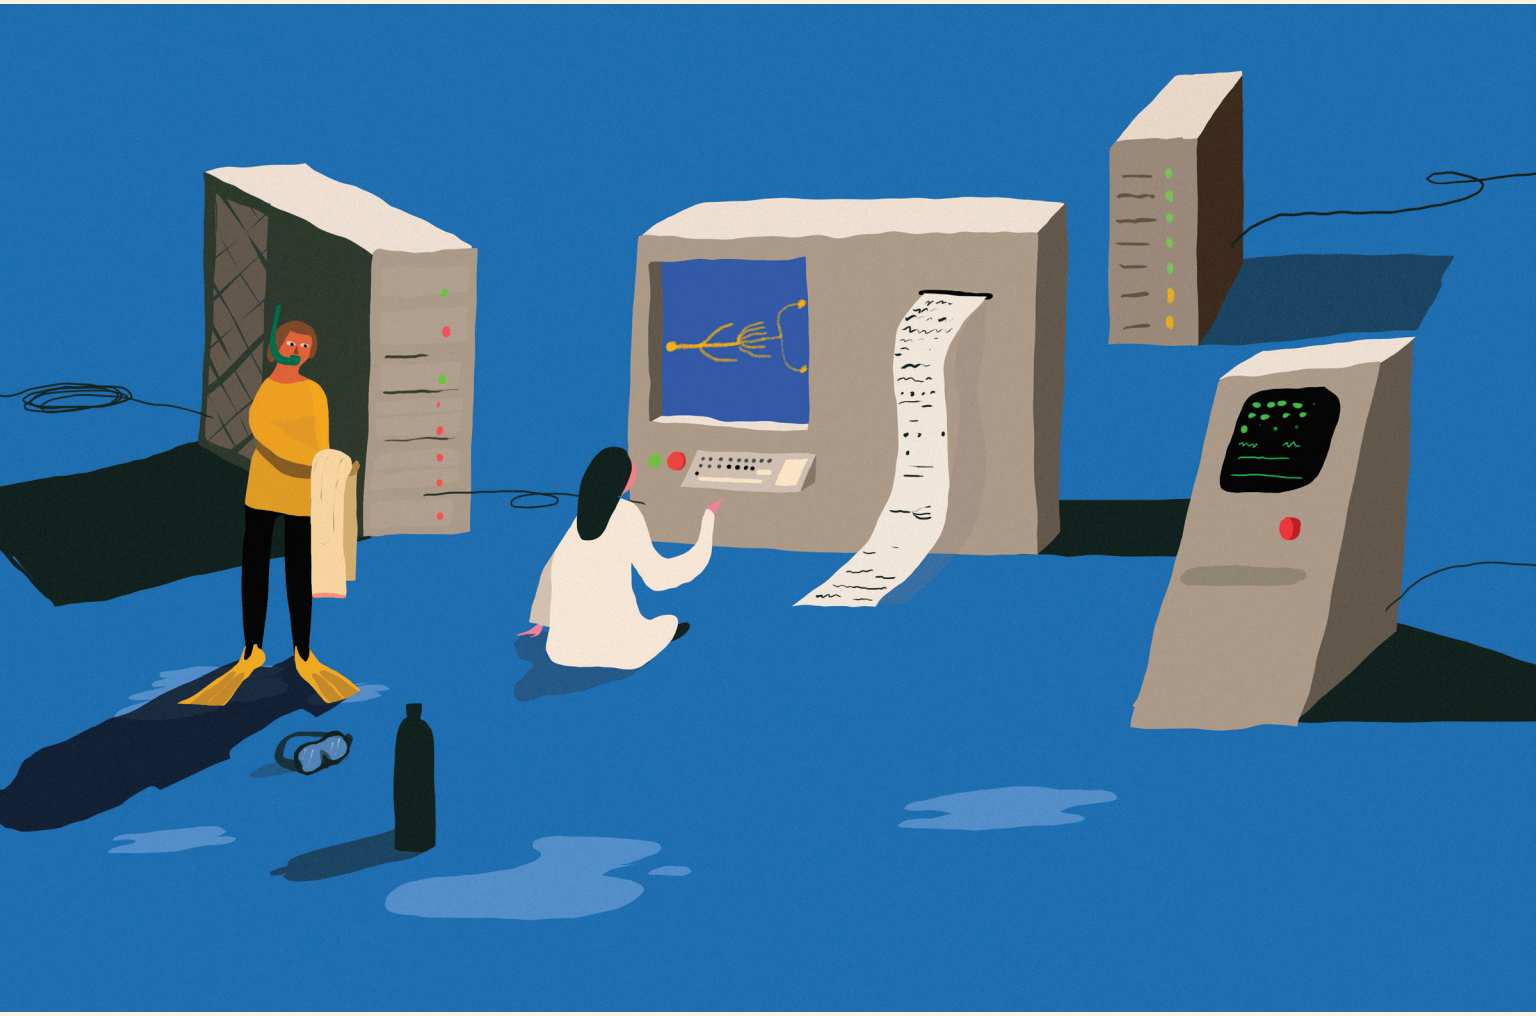

ILLUSTRATION BY DANIEL BARRETO

Our lab adopted practices to make data- and code-sharing feel safer, including having the coding equivalent of a dive buddy.

It is difficult to find someone in computational research who will publicly say they don't support open-science practices. By increasing transparency, open science accelerates progress and enhances equity. Sharing data and code makes the most efficient use of data collected from volunteer research participants, maximizes the return on public investments and helps investigators at under-resourced institutions.

Open science also breeds careful science: Researchers who know that their work will be open for external review may be less likely to cut corners. At the same time, though—despite support for open science and all its concomitant benefits—fear of this scrutiny can deter even well-intentioned investigators from fully participating.

To overcome this hurdle, my lab has adopted practices to make data- and code-sharing feel safer. Notably, we created a reproducibility buddy system—inspired by the buddy system in scuba diving, in which divers pair up to monitor and help each other in case of trouble. The practice is time-consuming, but it has helped us catch mistakes early and has made lab members feel more comfortable sharing their work.

BY TED SATTERTHWAITE, MCLURE ASSOCIATE PROFESSOR IN PSYCHIATRY AND BEHAVIORAL RESEARCH, UNIVERSITY OF PENNSYLVANIA

"Trainees report that the buddy system feels like a welcome safety net."

The impetus to develop this system grew from my own experience as a new faculty member. My biggest hang-up in embracing open science was simple: fear. Computational neuroscience can be complicated—lots of data and lots of code, all with many potential failure points. Code written by professional software developers includes approximately 20 to 70 bugs per 1,000 lines of code; code written by a brilliant but inexperienced graduate student likely includes far more. In a computational project with a complex code base, it is reasonable to assume there are many undiscovered bugs or errors.

The consequences of finding an error after results are published can be serious. Shortly after I started my lab, I found an error in a paper after it was published in a high-profile journal. Even though we caught the error ourselves, I was mortified. And the process of finding the error, understanding what results it affected, getting all stakeholders on board and working with the editor to correct the record was astonishingly time-intensive and personally draining. I imagine the process would have been even more stressful had an outside party relying on open code reported the error. Along the way, I also received an implicit message from colleagues: One erratum, you can still probably get tenure. Two? I wouldn't bet on it. This is not a culture that drives one to embrace open science.

Over time, we have changed the way we work so that open science feels less scary. Perhaps the single most useful among these is our "reproducibility buddy system." This framework—which I am sure is not novel—occurred to me after scuba diving. Scuba diving is generally quite safe but requires significant equipment because of the inescapable reality that we cannot breathe under water. To reduce potentially dangerous errors, one always dives with a dive buddy, who is responsible for helping check equipment and assist during a dive if problems arise.

Compared with a basic scuba outing, most academic projects are far more complex. If we are required to have a buddy to dive, why don't we have one for computational science?

The reproducibility buddy, which my lab members lovingly abbreviate to "reproducibili-buddy," serves a roughly analogous role. At the start of a project, we identify a team member to reproduce the project. To align incentives, this person is almost always the second author of the published paper. The reproducibility buddy replicates the work at several key checkpoints, going over every line of code and making sure they can run the code and get the same results.

Importantly, this process begins early in a project; nothing is worse than finding an error after a set of results has been polished to a high gloss ahead of submission to a journal. For example, one of the first major checkpoints occurs before we generate any results—we first reproduce all the nitty-gritty steps required to aggregate the data. Later, we ask the buddy to try to reproduce the first main result that becomes the anchor for any subsequent manuscript. Ahead of submission, the first author comprehensively cleans and comments on all code and creates a wiki that provides an overview of how to use it; the buddy's task is to reproduce the primary results using only this documentation. Finally, this process is updated as the work evolves in the revision process.

Sadly, this all takes time. Writing clean, well-commented code that can be easily replicated takes longer than hacked-together one-offs. Even under ideal circumstances, performing the replication itself is labor intensive. As in so many other domains, there is a real speed-versus-accuracy trade-off. Furthermore, the system is far from foolproof; results can be "reproducible but wrong"—code that runs and returns the reported result but is based on flawed scientific logic or a misunderstanding of the output.

When we first began piloting the buddy system five years ago, I anticipated pushback from the team. To my surprise, the opposite has occurred. Trainees report that the buddy system feels like a welcome safety net. Although replicating someone else's work is time-intensive, it is also a great exercise in code review that helps both parties learn from each other. Perhaps the single most common reaction is relief: The buddy system allows everyone to sleep better at night, knowing that their results have been vetted. With time, I have repeatedly learned that when we fail to use the checkpoints built into this buddy system, we do so at our own peril: Errors that should have been caught early on are found later and are far more costly to correct.

Implementing specific practices for reproducibility has helped our lab members be less afraid and even love open computational science. However, it is not a replacement for systematic change. As detailed elsewhere, there are multiple opportunities to encourage open science at every stage of the scientific life cycle. Granting agencies can prioritize data-sharing and re-use, open code and replication of findings. Top journals could expand options for published data descriptors to ensure credit for sharing data. Editors could enforce standards for data- and code-sharing, allow for registered reports and encourage updated results. Academic appointment and promotion committees could de-emphasize numerical measures of productivity and instead reward open practices and results that replicate. Perhaps most ambitiously, academic journals could compensate peer reviewers, making it financially feasible for results to be independently replicated as part of the review process. Although there have been encouraging developments on many fronts in open science, changing the existing consensus is inevitably slow. For now, we find that—as in diving—doing open science with a buddy is both safer and more fun.

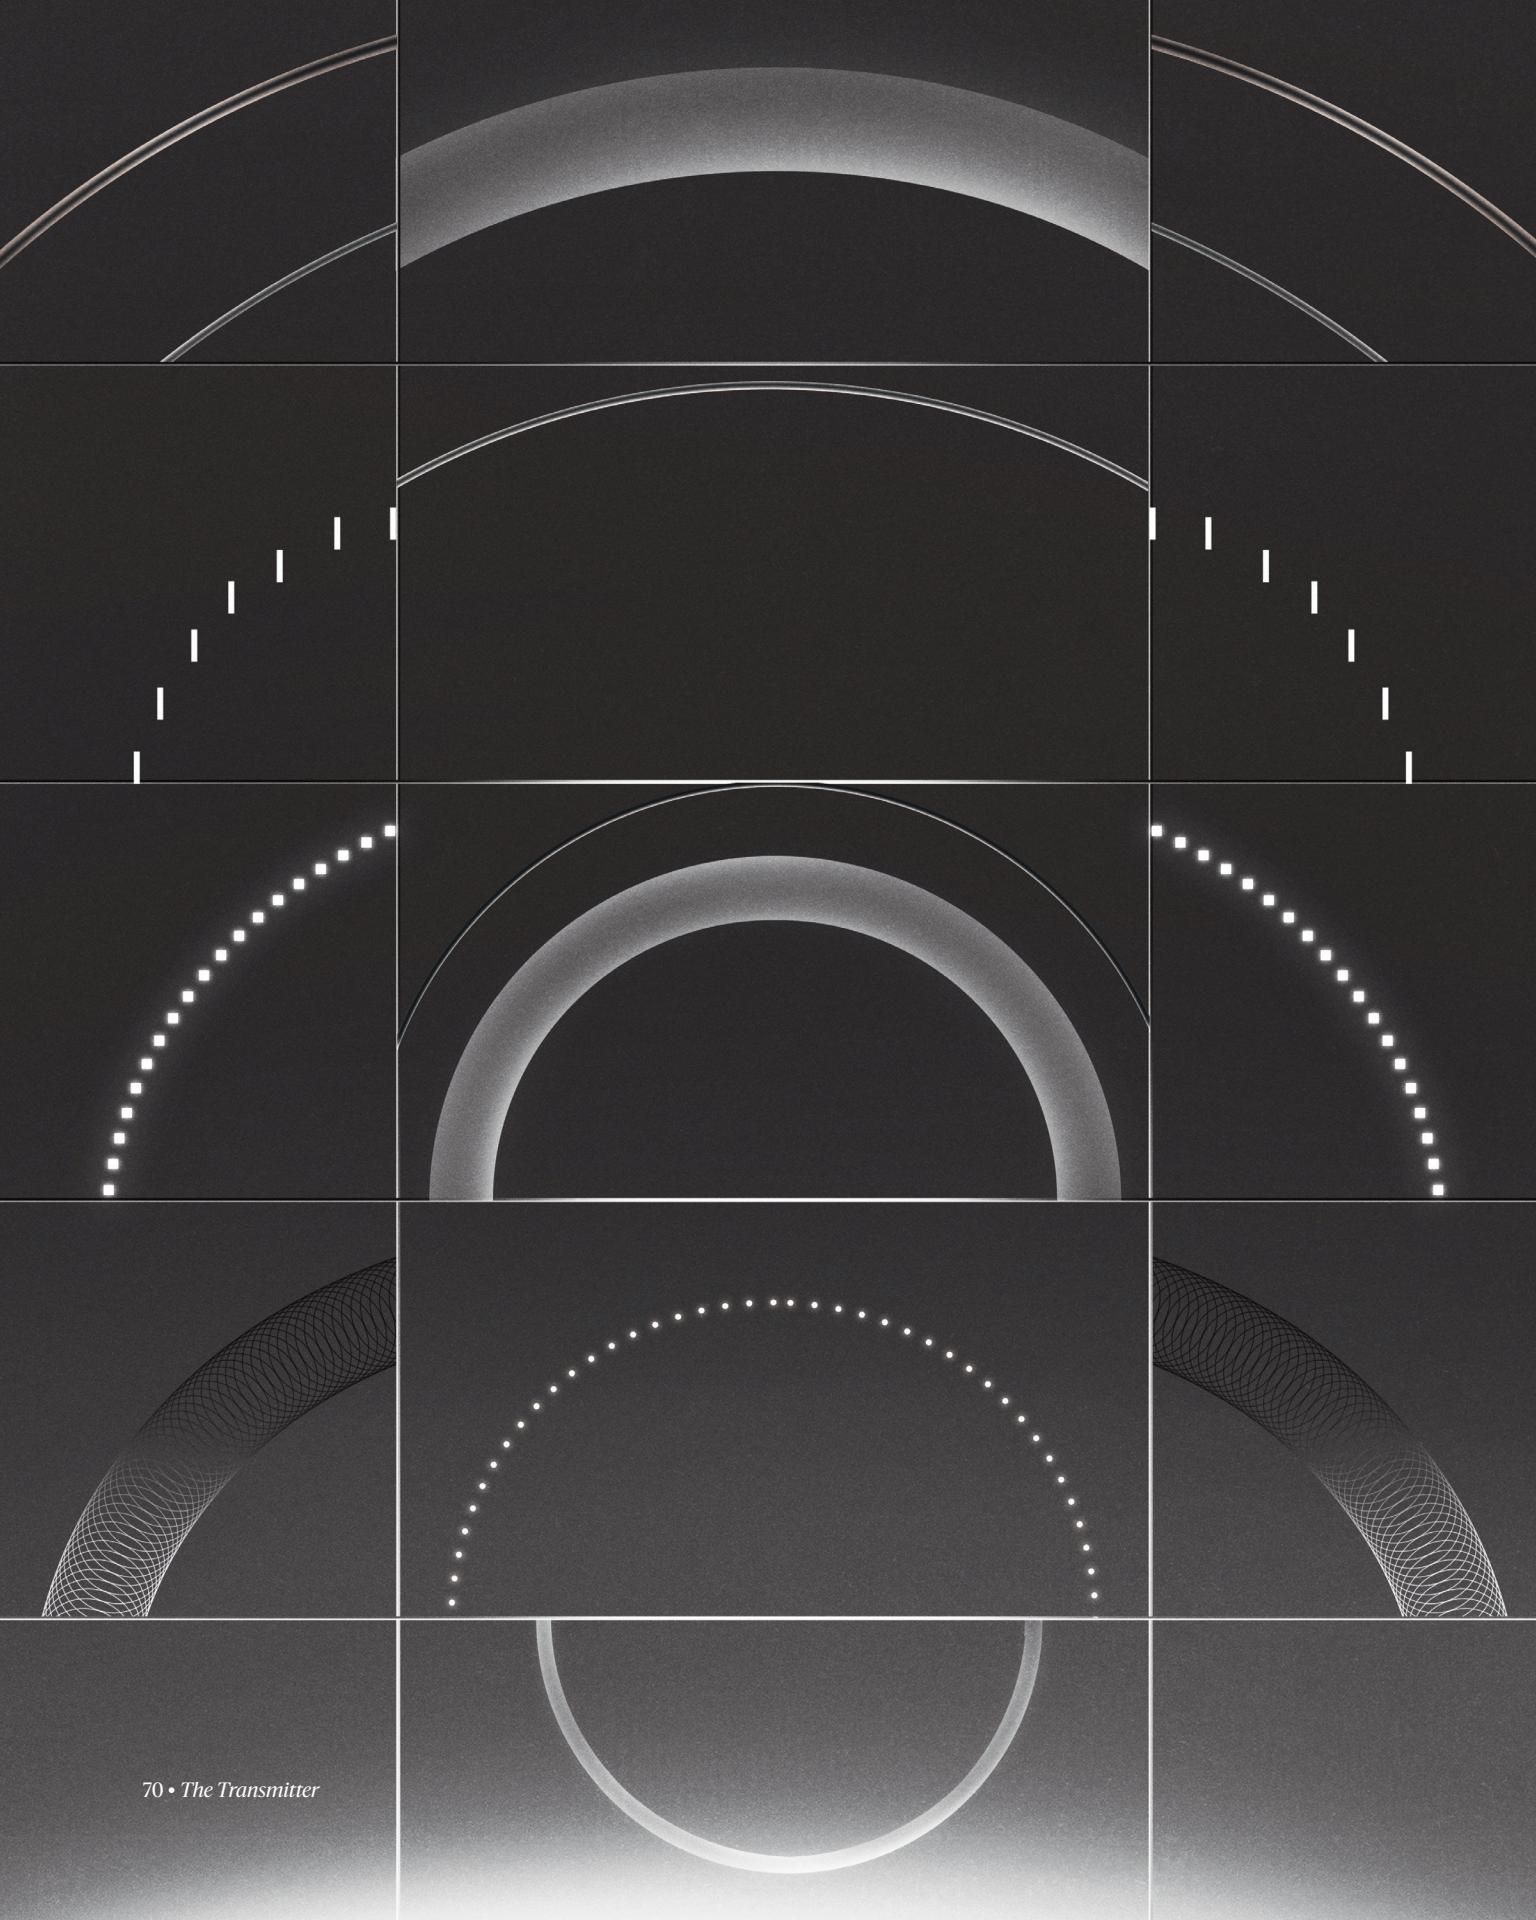

LOOKING AHEAD:

What do you think the field of neuroscience should prioritize for the next 10 to 20 years?

NICOLE RUST, PROFESSOR OF PSYCHOLOGY, UNIVERSITY OF PENNSYLVANIA; CONTRIBUTING EDITOR, *THE TRANSMITTER*

The field is increasingly embracing the notion that the brain is a "complex dynamical system" where causes lead to effects that feed back as causes—this happens through feedback loops within the brain and interactions between the brain and the environment. From ecology, engineering and other fields, we know that when complex dynamical systems go awry, they can be exceedingly difficult to restore. Tackling that challenge will be the key to developing treatments for the billions of people with brain conditions of nearly every type, from Parkinson's disease to psychosis.

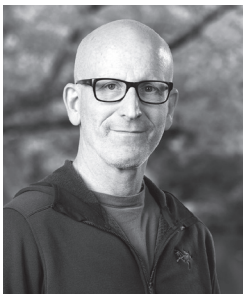

ANTHONY ZADOR, PROFESSOR OF BIOLOGY, COLD SPRING HARBOR LABORATORY; CONTRIBUTING EDITOR, *THE TRANSMITTER*

How do brains compute? Neuroscientists have learned a tremendous amount about the "parts" of the nervous system: the molecules and cells that make up the brain. What we still haven't figured out is how these parts work together to enable animals to outperform artificial intelligence on almost all tasks that require interaction with the real world: planning, sensorimotor interactions and balancing multiple goals—tasks that define an "embodied Turing test." Tremendous advances in computational and circuit neuroscience, as well as in AI, put these questions within our reach in the next decade or two.

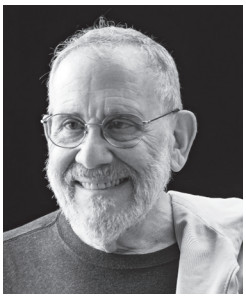

JOSHUA R. SANES, PROFESSOR OF MOLECULAR AND CELLULAR BIOLOGY, HARVARD UNIVERSITY; CONTRIBUTING EDITOR, *THE TRANSMITTER*

Mechanistic basic research on the human brain. We have learned enough from model systems over the past few decades that we can now apply these tools and insights to the human brain. Emerging or rapidly improving methods include organoids and assembloids, neuroimaging, extracranial stimulation (TMS, tDCS) and recording (EEG, MEG), single-unit recording over days or weeks in surgical patients, multiomics and spatial transcriptomics on postmortem tissue, transplants of humans neurons into mice (not chimeras), and brain-computer interfaces.

RUSSELL POLDRACK, COGNITIVE NEUROSCIENTIST, STANFORD UNIVERSITY; CONTRIBUTING EDITOR, *THE TRANSMITTER*

I think that a major priority for the next 10 to 20 years should be a shift from the current focus on data to a heavier focus on theory. The past two decades have seen the development of an amazing set of tools for the measurement and manipulation of neural systems, and with those tools has come an onslaught of data. Unfortunately, the initial optimism that more data would provide direct insight into brain function and structure has been dashed on the rocks of the immense complexity present in those data. I think that we need to balance the focus on increasingly sophisticated biological tools with more focus on the development of theories that can help us understand these massive data.

SHEENA JOSSELYN, SENIOR SCIENTIST, HOSPITAL FOR SICK CHILDREN; CONTRIBUTING EDITOR, *THE TRANSMITTER*

I'd like to take a step back because I am not convinced that it is a good idea to set priorities for neuroscience. I appreciate that several might answer this question by saying "translational studies, because we need to help those with brain disorders." I agree with the sentiment but would argue that for the vast number of brain disorders, we have little basic knowledge to translate. Instead, I think the most progress in understanding how the brain works (and therefore what to do when the brain isn't working so well) is built by encouraging researchers to follow their curiosity. We never know when and from where the next big potentially life-altering finding will be made. I would hate to see the potential ground-breaking discoveries of my colleagues limited by a top-down imposition of committee-defined priorities.

Defining

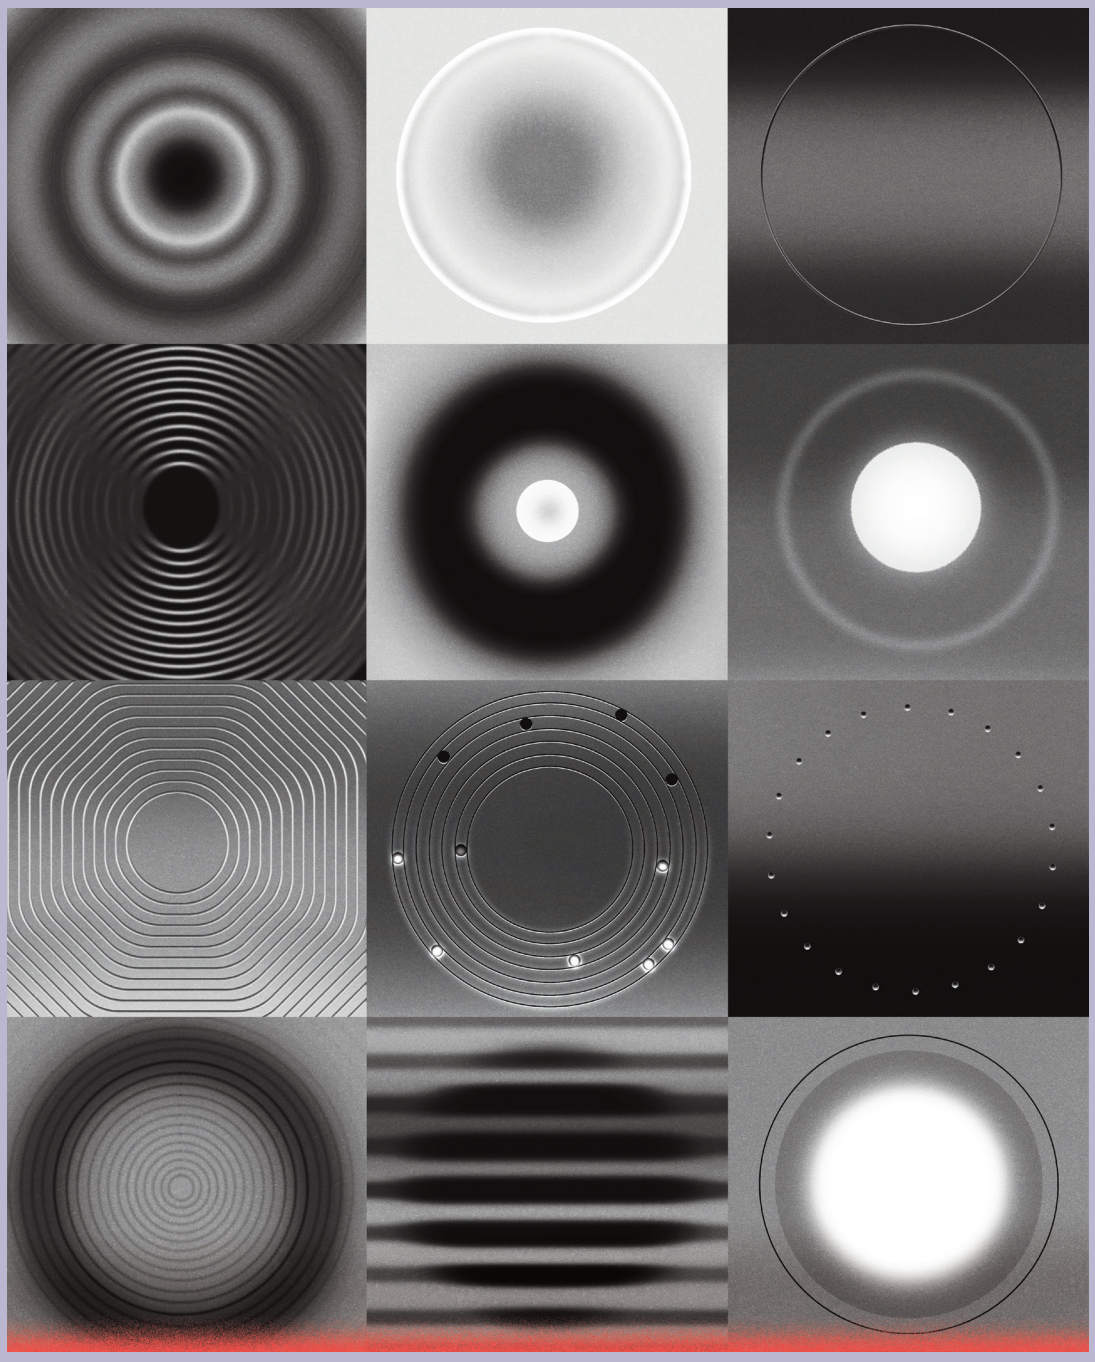

cell types

This series explores how new high-throughput technologies are changing the way we define brain cell types—and the challenges that remain.

Welcome to the second single-cell revolution: New high-throughput technologies are transforming how we define neurons

Cell census: The mouse brain has more than 5,000 cell types, categorized here based on their gene-expression patterns.

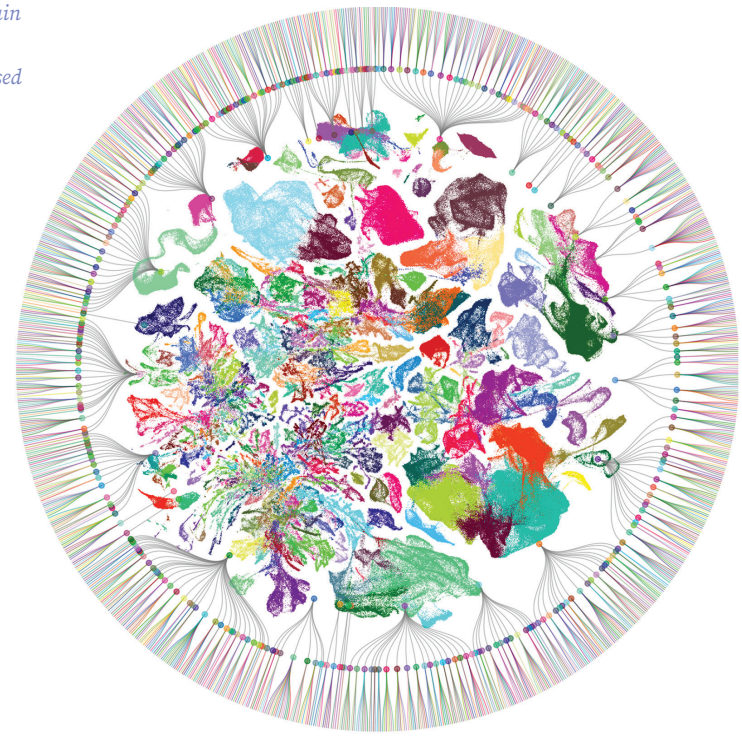

COURTESY OF THE ALLEN INSTITUTE

This ongoing essay series will explore questions these technologies raise, as well as opportunities they provide for understanding development, evolution and disease.

Although the brain has been an object of fascination for centuries, neurobiology as we know it originated with Santiago Ramon y Cajal's magnificent descriptions and drawings of neurons in the late 1800s. His work, which spanned nearly every part of the nervous system of dozens of invertebrate and vertebrate species, set the field's agenda for the next hundred years: the detailed analysis of single neurons, initially morphologically and later electrophysiologically. It was the first single-cell revolution.

By the end of the 20th century, adherence to this single-cell agenda—coupled with advances in molecular methods, including biochemistry, molecular biology, immunochemistry and transgenesis—had revealed the basics of neuronal structure, function and development. Despite these transformative advances, however, it became increasingly clear that to truly understand how the brain works—and how it fails in neurological and psychiatric diseases—would require new approaches. To make sense of circuit architecture and pinpoint cells that might be defective in brain diseases, for example, researchers would need to be able to study enough single neurons to classify them into types. That would require the capacity to analyze hundreds to thousands of single cells, preferably simultaneously.

BY JOSHUA R. SANES, PROFESSOR OF MOLECULAR AND CELLULAR BIOLOGY, HARVARD UNIVERSITY; CONTRIBUTING EDITOR, *THE TRANSMITTER*

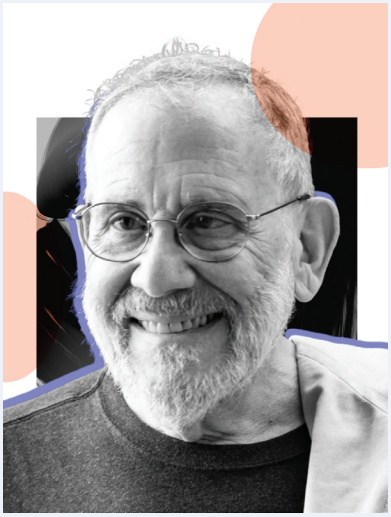

Brain maps: Researchers mapped transcriptionally defined cells across different regions of the brain.

That has finally become possible, with the invention over the past 20 years of a series of massively parallel high-throughput single-cell methods. Today, researchers can molecularly profile cells using high-throughput single-cell and single-nucleus RNA sequencing (scRNA-seq and snRNA-seq). They can monitor neuronal activity with multi-electrode "neuropixel" type probes and genetically encoded calcium and voltage indicators. They can visualize neuronal structure and connectivity with serial-section electron microscopy or optically with super-resolution and expansion microscopy. In each case, they can assay hundreds to thousands of cells in parallel sufficiently quickly and inexpensively (relatively speaking) to classify and characterize neurons and unravel neural circuitry more comprehensively than ever before.

In other words, we are now in the midst of a second single-cell revolution.

This new capacity to accurately classify cell types on a broad scale is revolutionizing how we study neural circuits, providing new insights into brain development and evolution, and opening new avenues for understanding brain diseases. But to fully realize its potential, the field still needs to grapple with a number of questions, such as what the most appropriate level to classify cells is and how closely different facets of cell-type data align. This introductory essay sets the stage for an ongoing series that will examine the uses of this technology for neurobiology, along with challenges that remain.

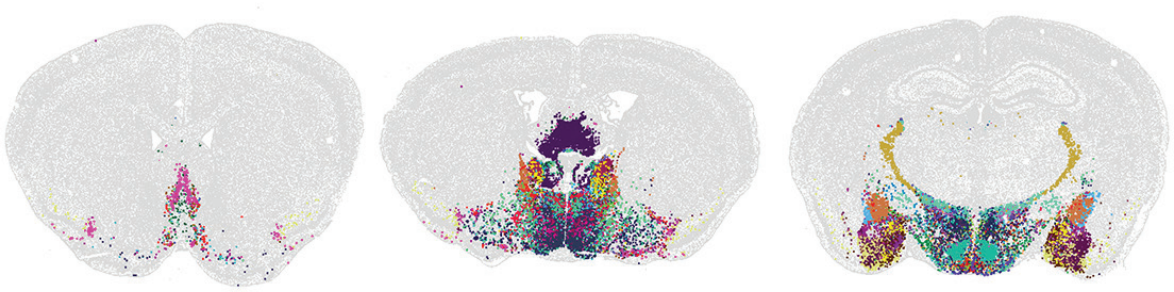

Of these new high-throughput methods, none has had a greater impact than sc/snRNA-seq. The method was invented independently and nearly simultaneously by three groups almost 10 years ago—Drop-seq by Evan Macosko and his colleagues, inDrop by Allon Klein and his colleagues, and 10X/GemCode by Grace Zheng and her colleagues. In all three platforms, mRNA from thousands of single cells or nuclei is captured and barcoded, then reverse-transcribed, amplified and sequenced in a single reaction. Cells with similar transcriptomes can be computationally grouped into candidate cell types. Newer methods provide similar results at an even lower cost than the original techniques. By now, researchers have profiled more than a billion single cells.

The first successful effort to use scRNA-seq to generate cell-type atlases of complex tissues used the mouse retina, a particularly accessible part of the brain; we now know that it contains some 130 neuronal types. Researchers have since applied the method to numerous other tissues and species. Over the past few months, Science and Nature have published special issues detailing the largest results from these efforts to date, including expansive atlases of the human and mouse brain.

The atlases, in turn, provide a foundation for addressing many important biological issues: What cell types are affected in neurological and psychiatric diseases? Where are the genes that predispose someone to or cause disease to be expressed? How do cell types that are resilient or vulnerable to insult differ? How do neural cell types diversify, differentiate and mature? How do activity-dependent and activity-independent factors influence these processes? Which cell types

"We are now in the midst of a second single-cell revolution. This new capacity to accurately classify cell types on a broad scale is revolutionizing how we study neural circuits."

are evolutionarily conserved and which arise to meet the needs of particular species?

Sc/snRNA-seq is still relatively new—it has been in wide use for just over five years—so studies have so far provided only partial answers to these questions. To fully address them, the field must overcome several technical and conceptual challenges, set forth below. Individual essays in this ongoing series will explore some of these specific questions in greater depth, including how cell-type data offer new insight into development and evolution, how to apply different computational methods for grouping cells, and the question of cell type versus state.

- **LOCALIZING CELLS:** Sc/snRNA-seq techniques begin by dissociating cells or nuclei, erasing information on a cell's location within the tissue. This is unfortunate because neurobiologists routinely rely on location—for example, to map connectivity between brain areas or target cells for recording. New "spatial transcriptomic" methods provide gene expression profiles of cells within tissues, using either multiplexed in situ hybridization to query a select set of genes (e.g., MERFISH) or RNA capture followed by sequencing (e.g., STARmap). Although these methods detect fewer genes than scRNA-seq does, they have become an indispensable adjunct to tissue profiling.

- **HARMONIZING CRITERIA:** Many researchers viewed the first scRNA-seq-derived atlases with skepticism because it was unclear whether cell types defined by molecular criteria corresponded to those that neurobiologists care about—structure, function and connectivity. In the retina, these different aspects of cell types align very well. But questions remain about other brain regions. Methods that record a single neuron's physiology, morphology and gene expression will help close this gap. Many to date are relatively low throughput, but combining calcium imaging with spatial transcriptomics holds great promise.

- **MULTIOMICS:** The transcriptome has proven useful for classifying cells but can't yet fully characterize them. RNA levels are imperfectly

correlated with protein levels and do not reliably distinguish among alternatively spliced isoforms. Moreover, they fail entirely to identify post-translational modifications, such as glycosylation, phosphorylation and many others, all of which are essential for neuronal function. This shortcoming has led to great interest in "multiomics," in which RNA-seq is combined with other methods. Researchers have successfully combined RNA-seq and ATAC-seq, which assays chromatin accessibility, a key epigenomic measure. But single-cell proteomic assays have not been optimized or widely used, particularly in combination with other methods.

- **GRANULARITY, TYPES AND STATES:** Researchers have not yet settled on the optimal resolution for grouping cells. At one extreme, there could be as many neuronal types as there are neurons. At the other, there could be very few types—for example, only sensory neurons, interneurons and projection neurons. Where is the sweet spot in between? This problem remains unsolved, with multiple computational models proposed to group cells into types, types into classes and so on. Perhaps more troubling is that initial transcriptome-based definitions of cell types made the implicit assumption that once animals reach adulthood, their transcriptomes are stable. Of course this is not true; neurons express different genes at different times, with major changes that depend on activity levels and patterns, hormones and more. Injury or disease lead to even greater changes. It remains challenging to distinguish these differences in cell state—meaning cells of the same type expressing different genes under different conditions—from differences in cell type.

Where do cell states end and cell types begin?

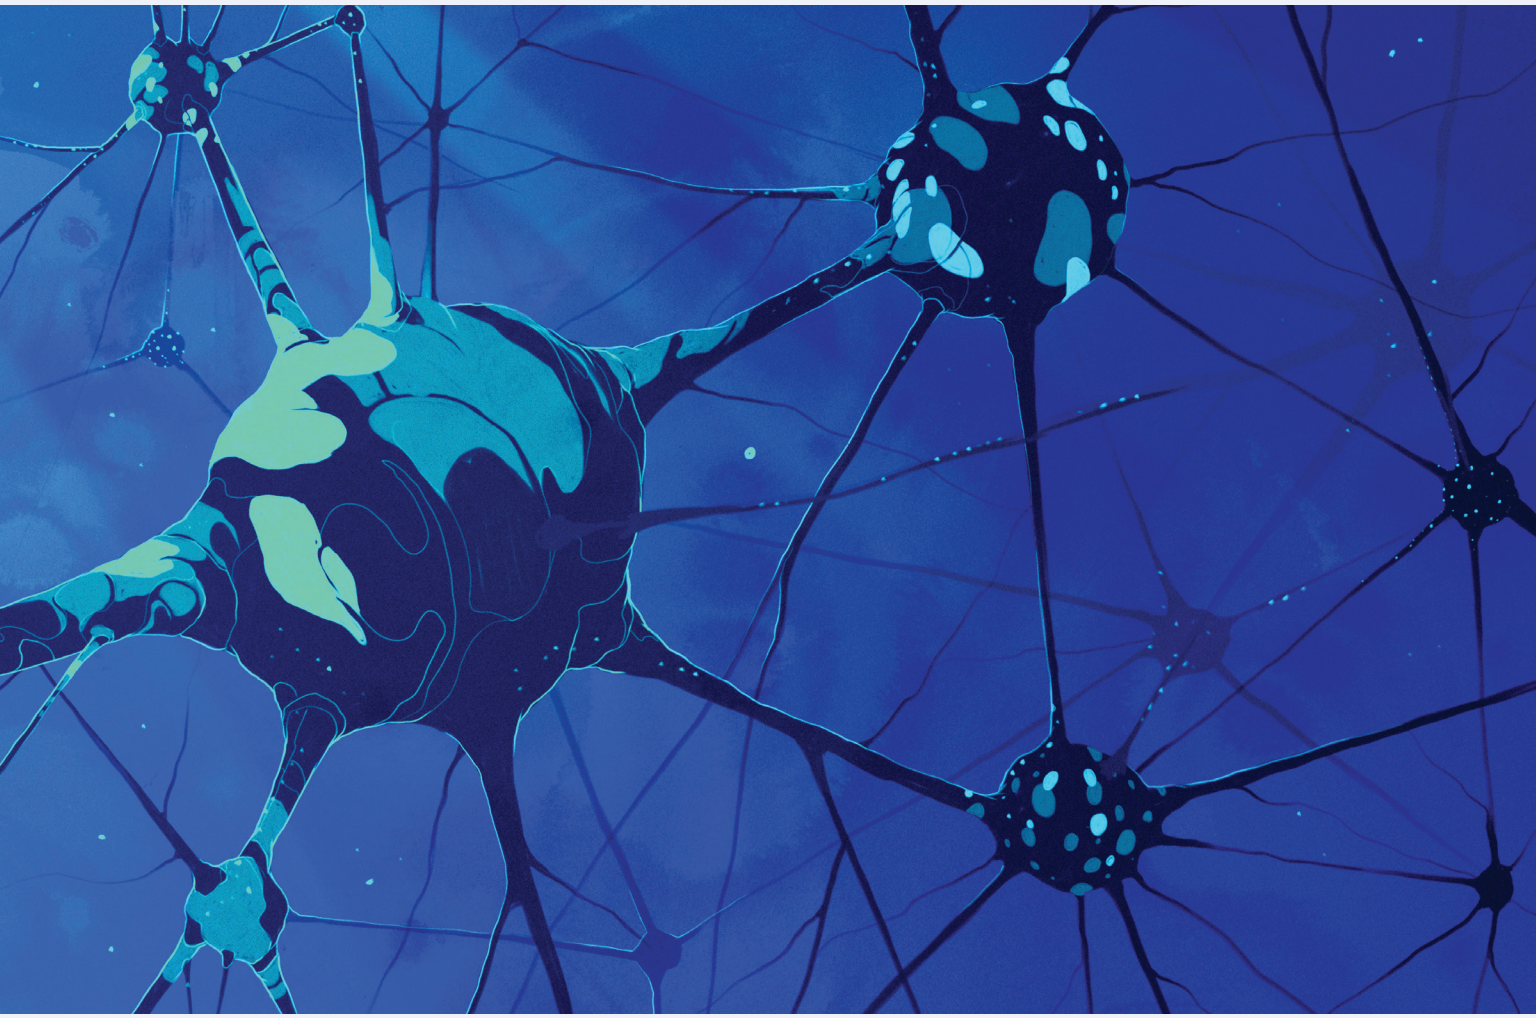

High-throughput transcriptomics offers powerful new methods for defining different types of brain cells. But we need to think more explicitly about how we use these data to distinguish a cell's permanent identity from its transient states.

"'Frogs are frogs and fish is fish, and that's that!' said the tadpole."

—LIONNI L. (1970). *FISH IS FISH*. PANTHEON BOOKS

In the children's story "Fish is Fish," a minnow and a tadpole celebrate their identity as fellow fish until the tadpole grows legs and hops out of the pond as a frog. The story reminds us that a collection of features observed at any one point in time is only a snapshot along the trajectory of any living thing, be it a cell or a more complex organism. Even if we introduced a subclassification of fish that includes both minnow and tadpole, without additional data these categories alone would not foretell that these creatures ultimately head in radically different directions from each other.

Neuroscientists risk falling into the same trap when it comes to cataloging the diversity of cell types that make up the brain. Single-cell and single-nucleus RNA sequencing (scRNA-seq and snRNA-seq) have revolutionized our ability to resolve the brain's heterogeneity—unsupervised algorithms can quickly classify cell types based only on the expression patterns of thousands of genes.

Despite its success, though, the reliance on transcriptomics to define cell types comes with intellectual hazards.

BY ANNE E. WEST, PROFESSOR OF NEUROBIOLOGY AND CELL BIOLOGY AND ASSOCIATE DIRECTOR OF THE MEDICAL SCIENTIST TRAINING PROGRAM, DUKE UNIVERSITY SCHOOL OF MEDICINE

"Given that gene-expression programs play a central role in defining both neuronal classification and cellular plasticity, how should we consider the question of where cell state ends and cell type begins?"

Like the fish story, defining cell types from transcriptomic snapshots assumes that a cell's gene expression is relatively fixed in time. Yet decades of evidence show that neurons undergo widespread and robust changes in their transcriptional programs in response to stimuli, including experience-driven neural activity. And these experience-induced transcriptional states can be quite persistent in contexts in which they mediate behavioral adaptability.

Given that gene-expression programs play a central role in defining both neuronal classification and cellular plasticity, how should we consider the question of where cell state ends and cell type begins?

Time-dependent transcriptional states are well understood in developmental biology. All the cells in an organism ultimately derive from the same source—a single fertilized egg—and all contain the same genomic DNA. Over the course of cell divisions and environmental exposures, progressive changes to the epigenome promote or restrict the expression of different genes, driving transcriptomic identities of distinct cell types.

Once established, epigenomic states can be remarkably persistent, such as the chromatin landscape that keeps one X chromosome permanently inactivated in each female cell. High-throughput, single-cell transcriptomic technologies rely on the stability of genome regulation to classify cells, and indeed, transcriptomic classifications of neurons overlap robustly with chromatin accessibility and DNA methylation patterns in single cells, supporting the premise of this classification strategy.

But the narrative of epigenomic inflexibility is inconsistent with current neuroplasticity research, which over the past 20 years has documented that numerous features of the epigenome can be modified by experience, even in terminally differentiated, post-mitotic neurons. If such fundamental mechanisms of genome regulation can change in a fate-committed cell, then researchers are left with an important question: Do differences in gene-expression programs always represent fixed cell types? Or could they also reflect transient cell states?

How we think about this question is shaped by both culture and technology. The microglia field offers an example of how naming conventions can influence how we interpret biological data. Historically, some groups referred to microglia in a way that implied static functional identities, akin to cell type, and others used language more reminiscent of cell state. A recent consensus paper on microglial nomenclature argues that the more static naming approach obscured an important aspect of microglial biology—that microglia transcriptomes are highly sensitive to the local environment. (Ironically, when describing neurons, the authors made the same mistake they had warned against, referring to neuronal transcriptomes as "fixed and terminally differentiated," ignoring their potential for plasticity.)

The techniques researchers use to analyze transcriptomics data, notably cell-clustering algorithms, also shape how we think about cell identity. These algorithms were explicitly designed to find discrete gene-expression programs that differ between tissues and represent cell types. They tend to overlook more subtle gene-expression programs in scRNA-seq data, those that vary over time or within areas and may reflect cell state. But clustering algorithms can sometimes detect cell state from scRNA-seq data as well—unpublished research suggests that they can identify neurons in a seizure state. Though this case represents a well-defined type of cell state, in which a single class of neurons is undergoing a precisely timed program of gene expression, it shows that the features we use to distinguish cell type from state may be less distinct than we think.

New computational approaches further support this idea. For example, Dylan Kotliar and his colleagues developed a mathematical model using matrix factorization that assumes cells can simultaneously express more than one gene transcription program, permitting cells to be assigned to more than one cluster. The researchers applied the model to snRNA-seq data from the visual cortex of mice that had been dark-adapted or exposed to light and showed they could identify activity-regulated transcriptional programs embedded both within and across cell-type identity clusters.

Studies that take a cell's precise location into account also support a more complex picture, identifying gene-expression programs that vary continuously across brain structures rather than in a discrete fashion. It remains to be resolved whether these gradient gene-expression programs should be conceptualized as subtypes of a cell type versus a single cell type that is varying its gene-expression state in response to its local environment. This question will become especially important as the field begins analyzing a major tranche of the data from the National Institutes of Health's BRAIN Initiative Cell Census Network (BICCN), published last October. These data will need to be placed into the context of brain circuits using spatial transcriptomics.

New experimental and computational methods will undoubtedly be essential in refining our understanding of cell types. But like the minnow and tadpole in the fish story, it will also be helpful to think outside the pond.

Building a brain: How does it generate its exquisite diversity of cells?

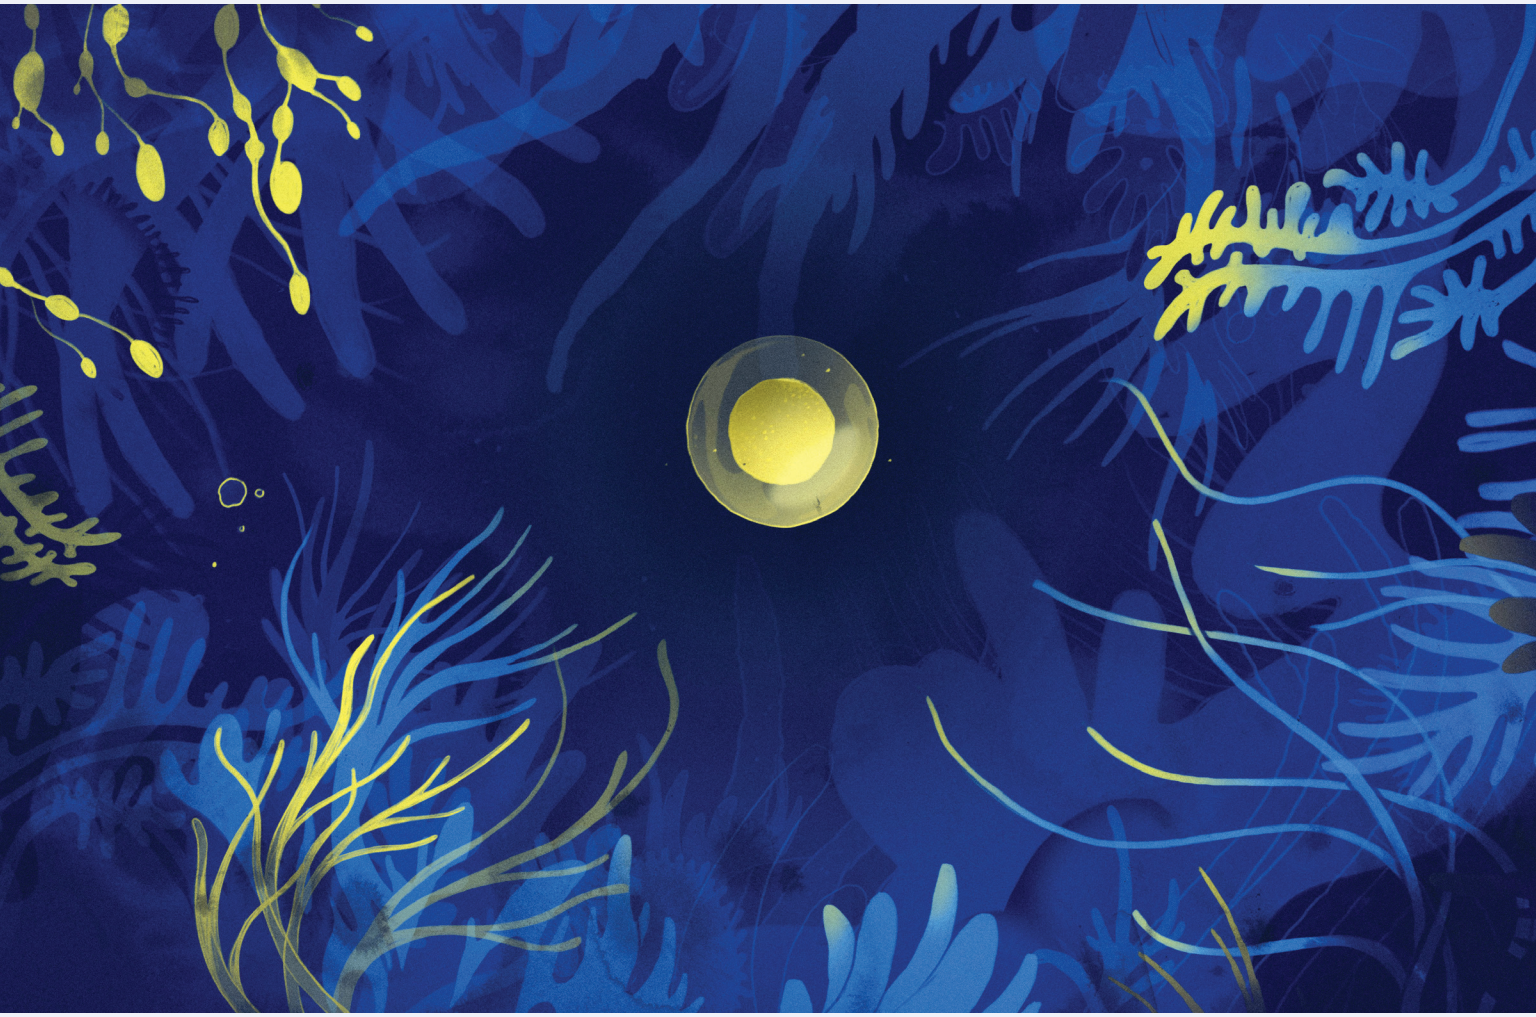

ILLUSTRATION BY SIMON PRADES

High-throughput technologies have revealed new insights into how the brain develops. But a truly comprehensive map of neurodevelopment requires further advances.

BY TOMASZ NOWAKOWSKI, ASSOCIATE PROFESSOR OF NEUROLOGICAL SURGERY, UNIVERSITY OF CALIFORNIA, SAN FRANCISCO; AND KARTHIK SHEKHAR, ASSISTANT PROFESSOR OF CHEMICAL AND BIOMOLECULAR ENGINEERING, UNIVERSITY OF CALIFORNIA, BERKELEY

Diverse neurons and their equally diverse circuits are the foundation of the brain's remarkable ability to process information, store memories, regulate behavior and enable conscious thought. High-throughput, single-cell profiling technologies have made it possible to classify these cells more comprehensively than ever before, offering a 360-degree view of the sheer magnitude of neural diversity in the mammalian brain. A recent effort to define the complete set of transcriptomic cell types in the adult whole mouse brain, for example, defined roughly 5,000 distinct cell types distributed across dozens of brain areas. This landmark accomplishment is a critical step toward integrating information about function and connectivity, and extending similar efforts to the adult human brain.

But this impressive gestalt conveys little, if any, information about how such diversity arises and develops in the first place. Single-cell atlases developed to date have been limited to a few points in time, focusing largely on the endpoint of neural development. How is this exquisite panoply of neurons generated and organized into precise and orderly circuits that last a lifetime? Providing the answer is the central task of developmental neuroscience. We want to understand the many transitions that unfold—where cells come from, the paths they take, and when terminal cell states emerge.

The comprehensive nature of single-cell technologies offers tremendous promise for defining cell types and reconstructing the trajectories of gene expression that underlie their differentiation. Initial efforts to apply these technologies to development, including in the prenatal human brain, hint at the insights these approaches can bring. Single-cell transcriptomics has helped map the diversity of neural progenitor cells, for example, most notably identifying progenitors that are expanded in humans, and their associated molecular adaptations. Further insights into development will require methods that reveal the specific history of every neuron type, including those that can more densely sample brain cells' trajectories over time and novel approaches for tracking fate transitions in individual cells. These discoveries will in turn help us to understand neurodevelopmental conditions, many of which are associated with genomic variation, and neurological disorders, such as brain tumors.

Studies of model vertebrates, such as rodents, frogs, fish and, most recently, primates, have generated important clues into how brain cells are born. The cardinal steps underlying brain development—neurogenesis, differentiation and the formation of initial synaptic connections—are largely regulated by intrinsic mechanisms encoded in the genome and do not require input from any extrinsic sensory experience. But external stimuli can extensively refine these circuits, so that each individual brain is custom-fitted to its unique external and internal worlds. Many genetically encoded pathways of brain development are remarkably conserved across the evolutionary tree, pointing to the ancient origins of the fundamental neurodevelopmental programs. Within this largely conserved scaffold, genetic variation abounds, underpinning species-specific features.

Despite these transformative discoveries, our understanding of the mechanisms that instruct neuronal diversification, maturation and wiring remains vastly incomplete. In the developing brain, neurons emerge from a limited pool of neural stem cells. Revealing the full picture of how neuronal types diversify and mature from their progenitors, we believe, will be significantly more challenging than simply classifying them. To understand the developmental history of any cell type, we need to know the evolving trajectories of the stem cells that contributed to that cell type, and to resolve the changes in molecular state, morphology, function, spatial location and connectivity along the way.

A single neural progenitor gives rise to many neuronal types through successive rounds of divisions, usually in a highly ordered, sequential manner. But neural progenitors come in diverse flavors. For example, mice have one type of radial glia, which produce one type of intermediate neuronal progenitor. Ferrets, monkeys and humans, in contrast, have three or more types of radial glia, and these may produce even more types of intermediate neuronal progenitors. We still have a long way to go to understand how these different cell types contribute to development.

"Current efforts to catalog brain cell types only sparsely sample the many stages of development."

Current efforts to catalog brain cell types only sparsely sample the many stages of development. Building a satisfying picture of this process demands much denser sampling—by roughly two orders of magnitude, both temporally and spatially—than what is currently underway. The scale and complexity of such efforts will be huge, relying on a comprehensive sampling of the developing brain. It will require the analysis of different animals at developmental time points that span all intermediate stages of cellular differentiation, ideally using methods that preserve information about cells' locations in the developing brain. Large-scale data integration and machine-learning approaches will also play an important role in this quest.

An alternative to this dense sampling is to look for hints of a cell's history in its epigenomic information, in particular its DNA methylation, which can harbor permanent marks of transcription factors' past activity. Or to employ tools that individually "mark" cells with unique sequences that "track" a cell from a point of origin, where the barcode is introduced, to a terminal position in the adult brain. These powerful approaches can reveal relationships among cells long after they depart from their original progenitor cell and position. Another way is to look at the DNA mutations that accrue in cells as they divide, creating a natural ticker tape to measure distances between cells in time. The complete lineage history is thus encoded within the genomes of individual cells, although it is challenging to comprehensively read out this information with currently available technologies.

In many cases, newly born neurons migrate long distances from their site of birth to populate different brain regions. Deciphering the developmental history of neuronal types will require tracing these migratory paths and pinpointing the molecular states of precursors as they traverse them. Postnatal neuronal migration is one of the major differences between mice and humans, and an area of active investigation. Using spatially resolved genomic approaches to profile the developing brain will be an integral component of these efforts in the near future.

In the first essay in this series, author Joshua Sanes described two single-cell revolutions made possible by technological advances that revealed features of neurons that were missed or absent in previous assays or approaches. We believe that generating a comprehensive map of neurodevelopment will be a third revolution.

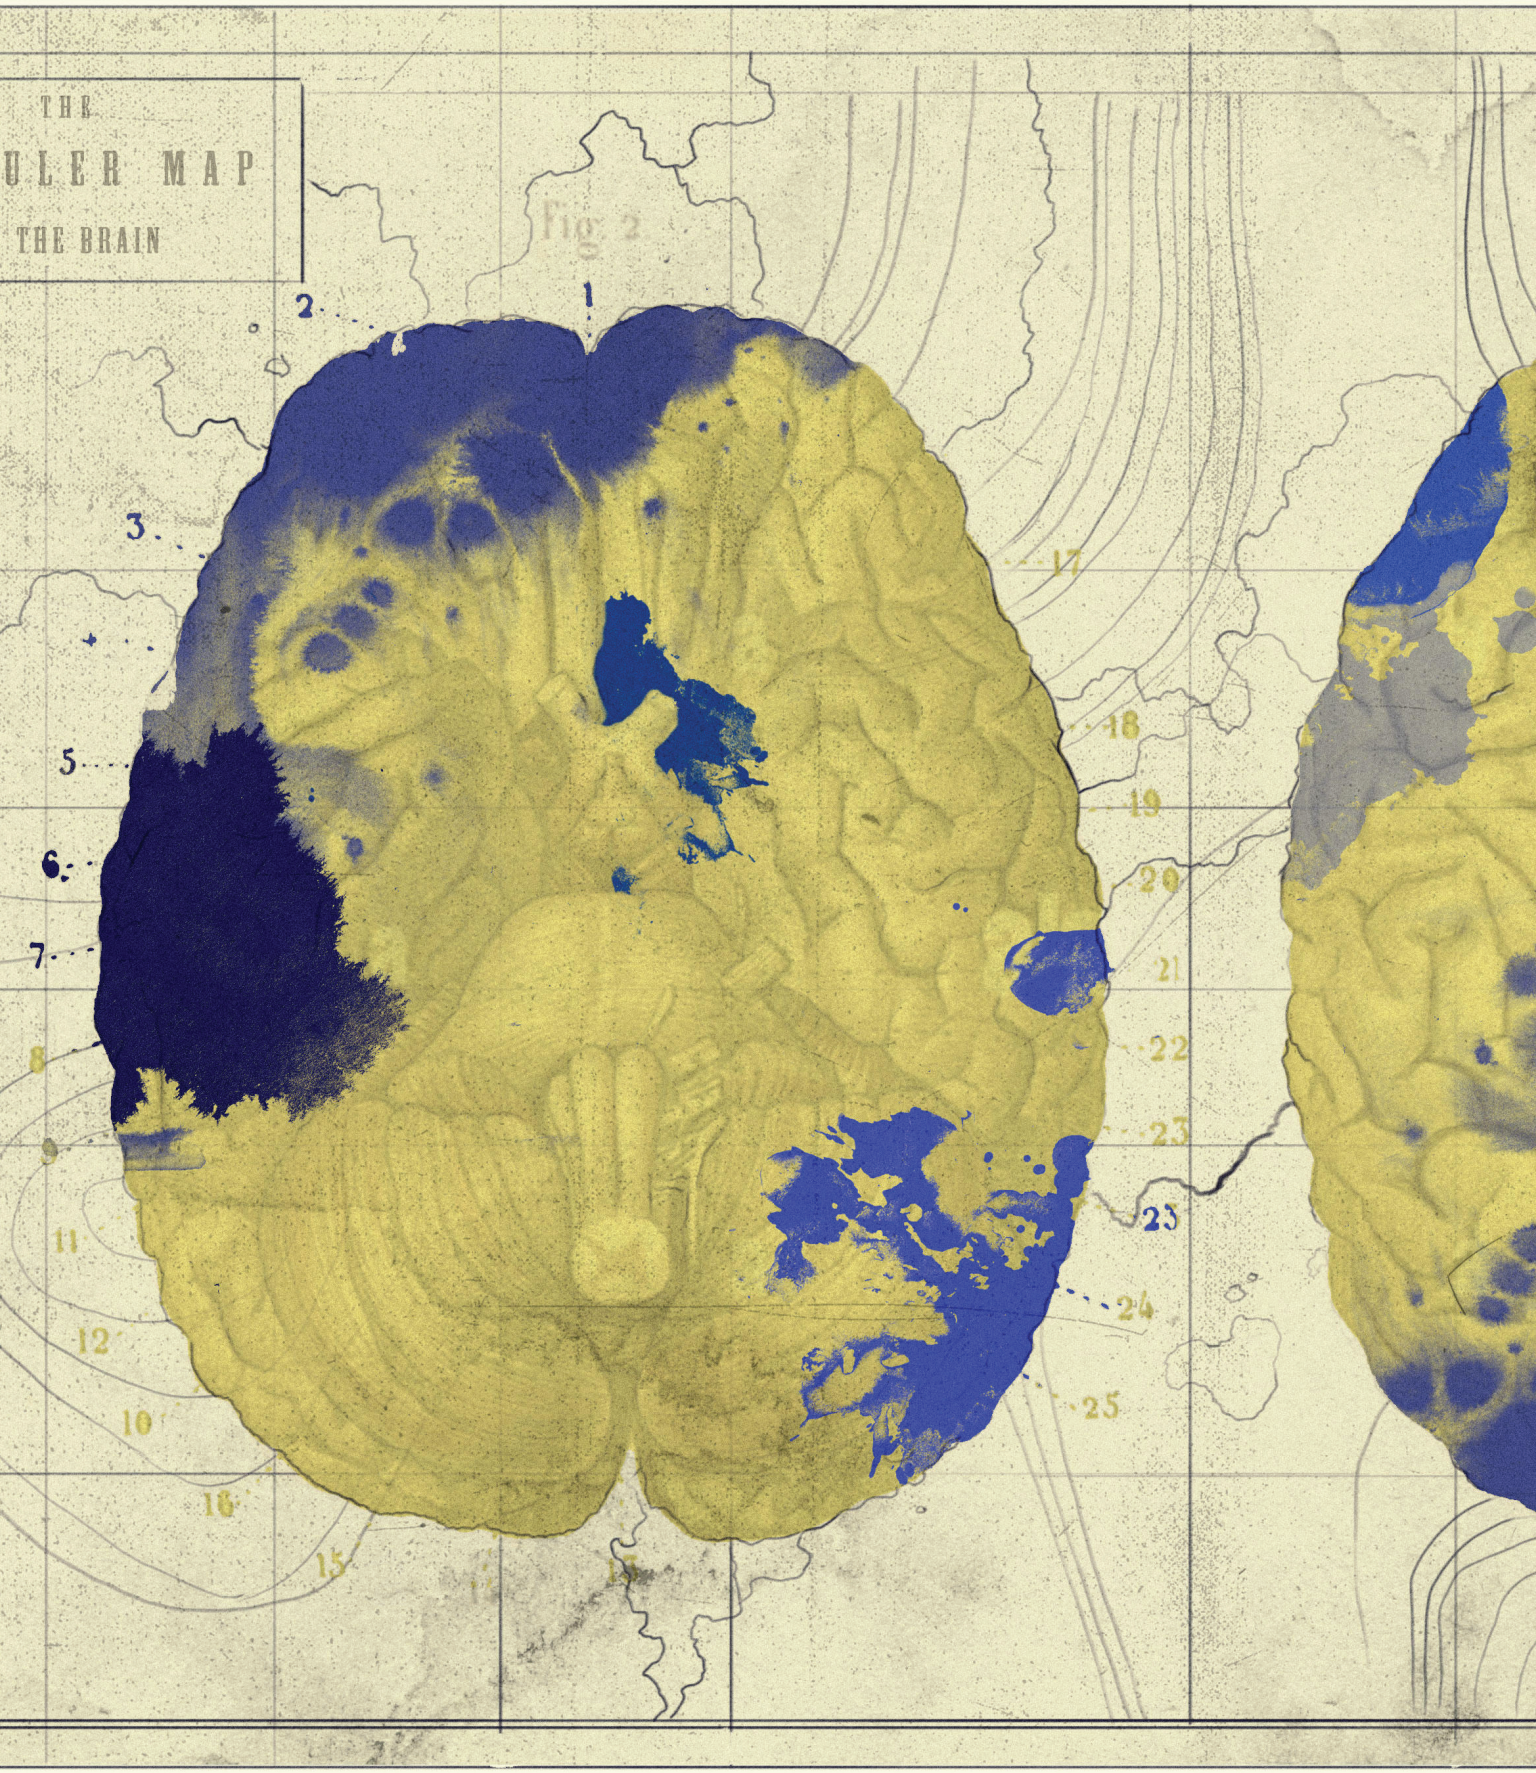
THE
ULER MAP
THE BRAIN
Fig 2
1
2
3
5
6
7
8
9
11
12
10
16
15
17
18
19
20
21
22
23
23
24
25

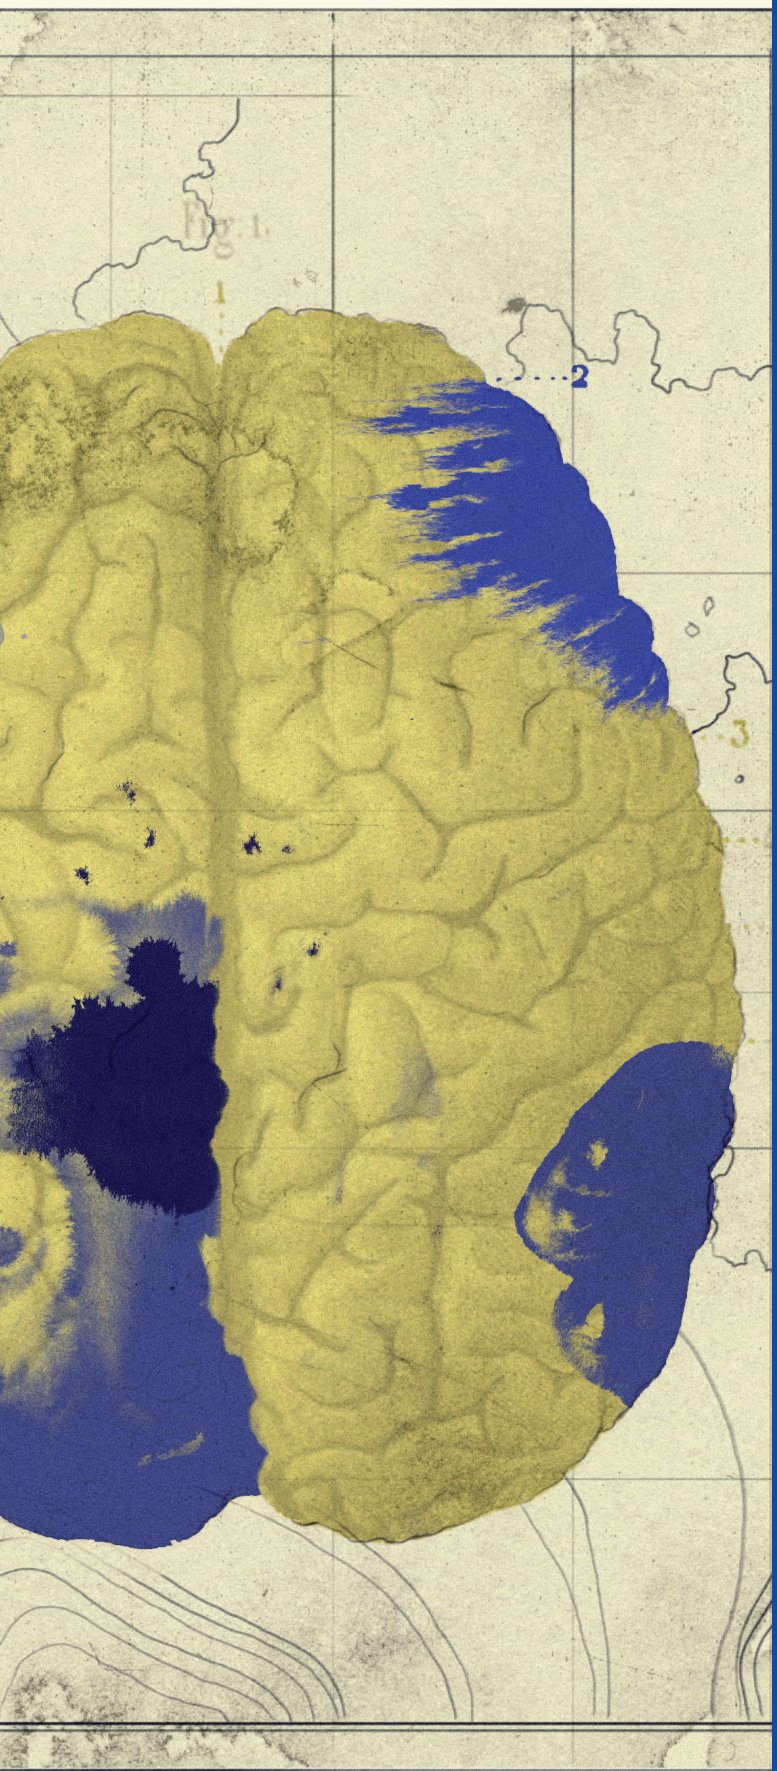

ILLUSTRATION BY SIMON PRADES

Knowledge graphs can help make sense of the flood of cell-type data

thetransmitter.org/defining-cell-types/knowledge-graphs-can-help-make-sense-of-the-flood-of-cell-type-data/

These tools, widely used in the technology industry, could provide a foundation for the study of brain circuits.

BY MICHAEL HAWRYLYCZ, INVESTIGATOR, ALLEN INSTITUTE FOR BRAIN SCIENCE

BIG PICTURE:

What, if anything, makes mood fundamentally different from memory?

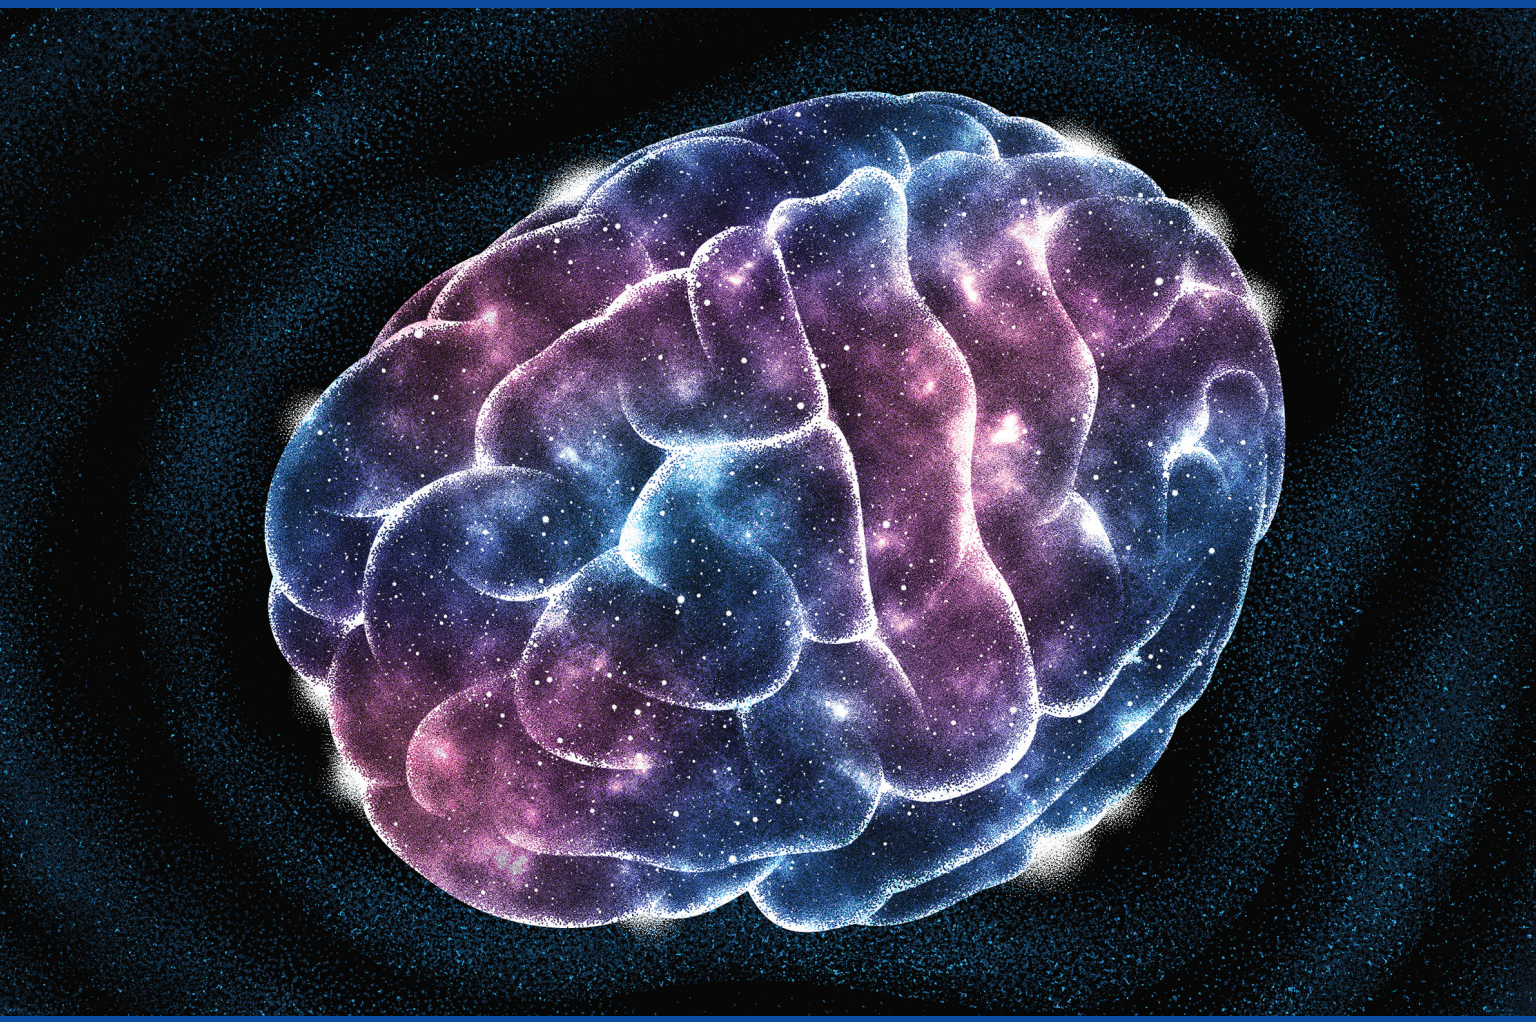

ILLUSTRATION BY KOUZOU SAKAI

To better understand mood disorders —and to develop more effective treatments—should we target the brain, the mind, the environment or all three?

We readily (and reasonably) accept that the causes of memory dysfunction, including Alzheimer's disease, reside in the brain. The same is true for many problems with seeing, hearing and motor control. We acknowledge that understanding how the brain supports these functions is important for developing treatments for their corresponding dysfunctions, including blindness, deafness and Parkinson's disease.

Applying the analogous assertion to mood—that understanding how the brain supports mood is crucial for developing more effective treatments for mood disorders, such as depression—is more controversial. For brain researchers unfamiliar with the controversy, it can be befuddling. You might hear, "Mental disorders are psychological, not biological," and wonder, what does that mean, exactly? Experts have diverse opinions on the matter, with paper titles ranging from "Brain disorders? Not really," to "Brain disorders? Precisely."

Even though a remarkable 21 percent of adults in the United States will experience a mood disorder at some point in their lives, we do not fully understand what causes them, and existing treatments do not work for everyone. How can we best move toward an impactful understanding of mood

BY NICOLE RUST, PROFESSOR OF PSYCHOLOGY, UNIVERSITY OF PENNSYLVANIA; CONTRIBUTING EDITOR, *THE TRANSMITTER*

and mood disorders, with the longer-term goal of helping these people? What, if anything, makes mood fundamentally different from, say, memory? The answer turns out to be complex and nuanced—here, I hope to unpack it. I also ask brain and mind researchers with diverse perspectives to chime in.

Among contemporary brain and mind researchers, I have yet to find any whose position is driven by the notion that some force in the universe beyond the brain, like a nonmaterial soul, gives rise to mood. Rather, the researchers generally agree that our brains mediate all mental function. If everyone agrees that both memory and mood disorders follow from things that happen in the brain, why would the former but not the latter qualify as "brain disorders"?

For many researchers, the debate centers on the level of focus: Does the brain or the mind cause fluctuations in mood? By analogy, consider Huntington's disease, a neurodegenerative disorder caused by a mutation in a single gene. We regard Huntington's as a genetically inherited brain disorder, not a subatomic particle disorder. The genetic mutation is, in fact, caused by a rearrangement of subatomic particles. But that's not a helpful way to think about the causes of Huntington's, insofar as our goal is to treat and cure it; it's simply the wrong level. Likewise, some researchers argue that the brain is not a useful level to think about the causes of mood and mood disorders insofar as we want to understand and treat them.

To elaborate a bit more, consider the phrase "insomnia causes fatigue." No one disputes it. But what do we mean by it? There's not a physical thing in the world, insomnia, that acts directly on another physical thing, fatigue. Rather, we mean that a lack of sleep causes the brain state that leads to the mind state of fatigue; in this case, insomnia and fatigue are not physical things but abstractions. If we want to help someone with insomnia, the most obvious way to intervene is not to determine the configuration of that person's ion channels or which of their brain areas are underactive, but rather to help them figure out how to get some sleep. By extension, some researchers propose that understanding mood at the psychological level, such as how a person experiences rewards, punishments and surprises relative to their expectations and beliefs, will have the greatest impact. Likewise, they argue that the most effective treatments will also be deployed at that level, in the form of behavioral interventions, such as cognitive behavioral therapy, mindfulness and other forms of psychotherapy. This is what researchers are referring to when they describe depression as psychological, not biological. By this logic, mood is different from at least some types of memory impairment, because behavioral interventions, such as memory training, cannot halt the progression of Alzheimer's disease.

An extension of this idea emphasizes the critical role that our environment plays in our mental health. We know that trauma and stress can trigger mood and other types of mental disorders, and that poverty and depression have a bidirectional, causal relationship. To that end, some argue that the most effective way to improve mental health will be through the social and envi-

ronmental interventions that we already know work. Likewise, they argue that research into environmental causes and interventions is as important, perhaps even more important, than brain research. You might be surprised by some proponents of this position, including the neuroscientists Peter Sterling and Michael Platt and the former director of the National Institute of Mental Health, Thomas Insel.

On the other end of the spectrum are those who argue that all mood disorders are brain disorders. Researchers who adopt this position acknowledge that, though environmental factors play a role, the same is true for people with memory impairment or even diabetes; in this view, memory is conceptually no different from mood. Advocates concede that, yes, of course, we should research the psychological and social factors that exacerbate these conditions, as we do for memory and diabetes research. But the bulk of our research efforts and dollars should focus on the biological phenomenon that mediates disease—in the case of diabetes, the pancreas and insulin; in the case of memory and mood disorders, the brain.

Some advocates of this position argue that because psychological variables are not physical—as in the insomnia "causes" fatigue example—it may be exceedingly difficult or impossible to establish cause and effect without investigating their physical correlates in the brain. Finally, advocates of this position point to evidence that brain- and body-based interventions, including antidepressants and exercise, as well as noninvasive and deep brain stimulation, work for at least some people with mood disorders. Given how little we know about mood in the brain, it only makes sense to do more research on mood in the brain to determine if we can improve upon the brain-based therapies that already exist.

Many psychiatrists adopt a more nuanced position, in which the best level to explain the causes of mood and to target interventions for mood disorders depends on many factors, which can vary from person to person. For mood disorders, both medicine and behavioral therapies, most notably cognitive behavioral therapy, can be effective. By contrast, behavioral interventions are not effective for memory disorders to the same degree. And the most effective treatment strategy will differ for different people. Many who adopt this position advocate for more research at both the biological and psychological levels, and the interaction between the two, emphasizing the need to figure out how to predict the treatments that will work best for specific people.

A final, more up-and-coming position draws from notions in physics, emphasizing the need to acknowledge that mood is an emergent property of a complex system. Proponents of this idea point to the innumerable feedback loops within and among the brain, the mind and the environment, and their parallels to other complex dynamical systems, such as the weather and ecosystems. They argue that because complex systems cannot be understood by deconstructing them into their parts, mood must be understood holistically and simultaneously across levels, applying what some call a biopsychosocial approach. Here, the gist is that interactions between the brain, mind and environment must be investigated simultaneously to understand mood and treat mood disorders. In complex systems with emergent properties, the properties of

higher levels cannot be inferred from lower ones. Consequently, if mood is an emergent property of the brain, it cannot be understood by studying the brain alone. Similar ideas exist for memory, though they are less pervasive.

To better understand the different ways brain and mind researchers are thinking about this topic, I asked them:

What do you see as the most effective path forward for mood research? How does it compare with the path forward for memory?

I was struck by the diversity of these viewpoints, all rational but also quite different. All five of the perspectives I've just described (setting aside the sixth, a nonmaterial soul) are reflected in these responses.

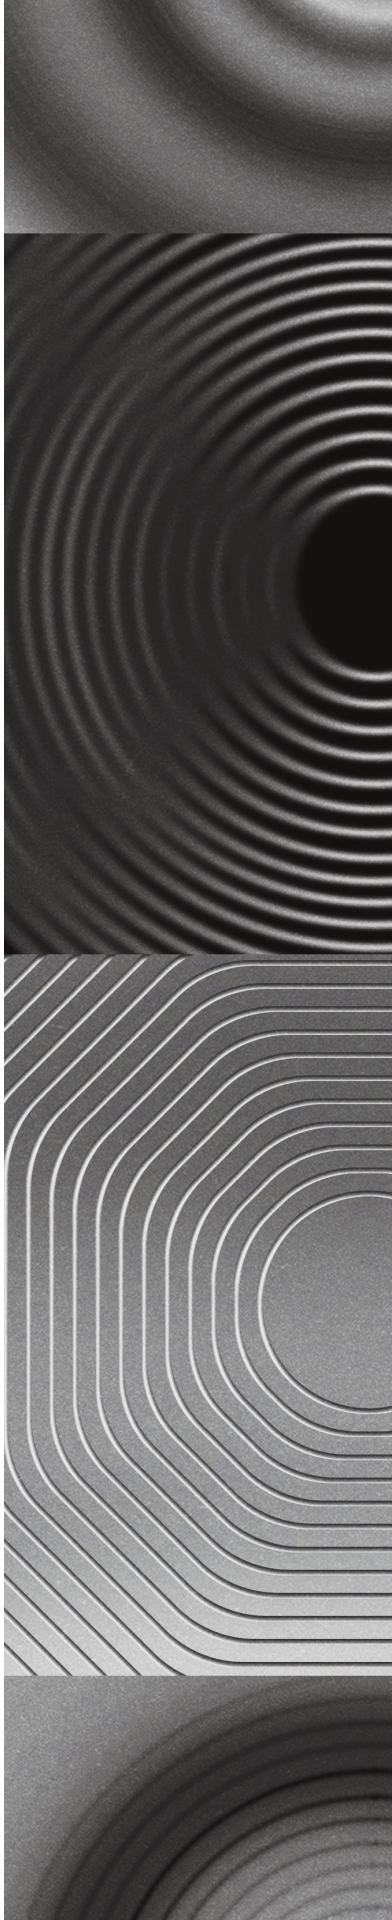

What do researchers think?

AWAIS AFTAB, CASE WESTERN RESERVE UNIVERSITY

Whether mood disorders (and mental disorders generally) are "brain disorders" is an interesting philosophical question. Unsurprisingly, the answer depends on how we understand the notion of a brain disorder. It is important to recognize that this question is motivated by daunting epistemic challenges in brain-behavior research. Neuroscientists have been fortunate that classic memory disorders have turned out to be cases of cellular neurology gone awry, with downstream effects on brain circuits and cognitive domains. Mood disorders, in contrast, appear largely to involve brain-body-environment interactions locked into dysfunctional patterns. Historically, neuroscience has lacked the tools to meaningfully articulate such higher-order dynamic interactions, and its current abilities to do so are fairly rudimentary; psychological concepts describe these interactions, albeit in a manner that is often idiosyncratic and socially constructed. Concepts such as narcissism, shame, projective identification—so familiar to psychodynamic clinicians—resist easy articulation in neurological terms. Maybe one day neuroscience will develop the language to adequately describe and explain these higher-order phenomena, but it is not there yet. Neuroscience requires periodic reminders that its current methods are too crude, that it should not be so arrogant as to think it can make psychology and the social sciences redundant. Mood disorders are multi-level phenomena, not categorizable simply as problems of "mental software," nor as problems of "neurological hardware." Their scientific understanding requires the "piecemeal integration" of multiple scientific disciplines and a great deal of epistemic humility.

AUSTIN COLEY, UNIVERSITY OF CALIFORNIA, LOS ANGELES

Investigating the nuances of mental health disorders, particularly in mood disorder research, holds immense promise for both basic science and clinical applications. Moving away from rigid diagnostic categories toward a spectrum-based approach can revolutionize our understanding and treatment of these conditions. This shift allows for a more precise diagnosis and personalized treatment plan tailored to individual needs, recognizing that not everyone responds uniformly to the same interventions. By closely tracking behavioral patterns alongside neural activity over time, we can uncover subtle yet significant changes at a granular level. This is the scientific basis of my research program. This comprehensive approach will enable us to fine-tune treatments and potentially offer preemptive measures for people predisposed to mood disorders. Also, examining the dynamic interaction between mood fluctuations and cognitive function, such as memory formation, is essential. Those with major depressive disorder commonly experience cognitive decline, making it imperative to elucidate the relationship between depressed mood and memory consolidation. By doing so, we can gain valuable insights into the underlying mechanisms driving these conditions, thereby leading to more effective interventions and improved patient outcomes.

EIKO FRIED, LEIDEN UNIVERSITY

For me, inroads for mood research in the mental health sciences will come from better understanding the systems these states play a role in. A person's mood states—happy, angry, anxious—are experiences that serve as both causes and effects. These states influence each other, but they also influence a person's thoughts and behaviors, and the environment. What I experience when I receive a bad grade in school or kiss someone influences—and is influenced by—my mood states. The systems theory of mental disorders posits that these networks of mood states and related features can form healthy attractor states: Reciprocal interactions and feedback loops keep many of us in a healthy state. An episode of major depression, then, is an alternative stable state in this system. If we take this perspective seriously, the next steps are

to map out these systems, by combining dynamic data, collected via smartphones, for example, with appropriate statistical models to study how these systems differ across people, how stable they are over time, and if systems change during transitions. We carry out some of this work in the WARN-D lab.

STEVEN HYMAN, BROAD INSTITUTE

The most certain path to better treatments for mood disorders depends on insight into their mechanisms—and thus into brain biology. Whether the onset of a depressive episode follows an adverse experience such as loss of a job or seems to come from nowhere, it is grounded in neural mechanisms, as are the ultimate treatment targets, regardless of how treatment is delivered—via psychotherapy, medicine or electrical stimulation. Such assertions do not represent reckless reductionism. Brains support the fitness of all free-living animals, including ourselves, because they are remarkably powerful integrators. To select actions and regulate our physiology in response to threats, rewards or complex social interactions, the brain must synthesize information that comes from each person's DNA, prior developmental events and their bodies, with current sensory inputs that refine its predictions of what is happening in the present and what will happen in the future. To discover better treatment, we must understand how such mechanisms go awry. Some who accept the brain as the substrate of thought, emotion and behavior—nature's lesion "experiments" are quite convincing—but nonetheless devalue the brain's importance, often hold an erroneous model that analogizes the brain to a digital computer that could run any software. They construe the neural hardware as a dumb machine programmed by the important stuff, thought and experience. But the brain is nothing like a digital computer. There is no software independent of biology. Even as neural circuits are computing the current moment and predicting the next, the brain's structure, including molecules, dendrites and synapses, is changed forever. It is the biology and plasticity of the brain itself that matter for all experience and action and all psychiatric illness. There is nothing left over.

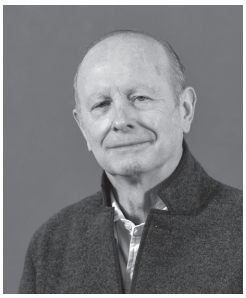

JOSEPH LEDOUX, NEW YORK UNIVERSITY

The ability of mental health professionals to deliver treatments that effectively relieve mental anguish pales next to the success that other areas of medicine have achieved. I think part of the problem is that the medical model may not be as applicable to mental disorders, or at least not to the mental part of such disorders. Medicinal treatments are especially useful in altering behavioral and physiological symptoms. But mental anguish is subjective, conscious suffering. That does not mean that conscious experience is not physical. It just means that it involves higher-order circuits that operate independently of those that control behavioral and physiological responses. It is not surprising that medications that alter behavioral and physiological responses in animals do not do much to relieve mental anguish in humans. Possibly medications developed in human studies would do a better job. But only if the researchers accept that changing a person's behavior or physiology is not going to do the trick. The target has to be subjective experience itself. The core problem of mental anguish might be best treated with psychosocial approaches. Medications can also have a role in helping to ease process, but not as a cure or primary treatment.

LISA MONTEGGIA, VANDERBILT UNIVERSITY

My lab studies synaptic plasticity processes in memory, as well as in the treatment of mood disorders. I think of memory and mood as manifestations of brain function that rely on processes from neurons and their synaptic contents. These two distinct processes may be embedded within the same neuronal synaptic network yet have distinct features. Forming memories requires encoding precise bits of information in our synaptic networks. Precision, long-term storage and retrievability are essential features of memory. Mood, in contrast, is a global regulator of synaptic networks that does not necessarily influence the specific information encoded but rather determines how and when it is encoded and retrieved, and what actions we choose to take in response to this information. In essence, memory is the information, whereas mood is how we use and react to this information. Therefore, these two processes may arise from the same synaptic networks in the brain but rely on their distinct abilities to process information.

RUSSELL POLDRACK, STANFORD UNIVERSITY

Memory and mood are clearly similar in the sense that they are defined by our subjective experience and can have powerful effects on how we think, feel and behave. But I see at least three ways in which they differ, which are directly relevant to how we can most successfully study them. First, memory is (at least often) tethered to the objective world, whereas mood is wholly subjective; whether I went to Astroworld as a child is a matter of fact, but whether I am happy today is simply a matter of my subjective experience of happiness. This external verifiability of memory has made it much more tractable to study, and particularly to develop models in nonhuman animals. Second, memory poses a much clearer set of computational problems, such as recognition (differentiating previously experienced stimuli from similar but novel stimuli) and recall (retrieving details about previous experiences). This has led to the development of successful computational models of memory that have made it possible to computationally characterize neural circuits for memory. The computational problem that mood solves, and whether it even makes sense to analyze mood in computational terms, are much less clear. Finally, with memory, we may have simply gotten lucky with regard to functional anatomy. The fact that limited lesions to a small set of brain structures in the medial temporal lobe can result in punctate memory disorders has provided a platform for subsequent analysis and modeling of those neural circuits. There does not appear to be any similar high-value target for mood. Instead, it seems that mood disorders can result from lesions across a widespread set of brain regions. The "localize and decompose" strategy may have worked well for memory, but mood may require a different strategy, such as one grounded in complex systems theory.

LAUREN ROSS, UNIVERSITY OF CALIFORNIA, IRVINE

Memory and mood both have genuine causes within and outside of the brain, but they also differ in important ways. In particular, it is worth considering whether mood states have more psychological causes and characterizations (in contrast to biological and neurobiological causes), especially when compared with memory. One main reason for this is evidence that moods can be successfully and strongly influenced by "self-interventions," such as when people alter their own thoughts and

emotions as a way of changing their behavior. Cognitive behavior therapy, talk therapy, and mindfulness and meditation all fall into this category. Though memory can also be influenced by psychological factors, it seems somewhat less subject to self-intervention. And memory has other features that make it easier to connect to neurobiological-level causes—memory is more straightforwardly conceived of as a coherent, functional system, whereas "mood" is often an umbrella term for disparate states, such as depression, anxiety and fear, many of which gain scientific attention as dysfunction or disease, rather than functional systems. Further exploring these similarities and differences is important for supporting advances in both.

ROBB RUTLEDGE, YALE UNIVERSITY

Our understanding of mood at a computational and neurobiological level lags behind our understanding of memory. I think this is largely because of mood's subjective nature. We can ask an animal which picture they have seen before, but we can't easily ask an animal if they are happy. Can we just use neural and physiological measurements instead? Not quite. A person can be wrong about whether they have seen a picture before, but they can't really be wrong about whether they are happy, no matter what our measurements say. That doesn't mean we can't ask people what they are feeling and use mathematical models to predict what they will say. This approach has shown that happiness depends on whether you are doing better than expected recently, and that sometimes learning matters more for happiness than reward (you might have a hobby where this finding applies). It turns out that happiness can sometimes go up and down in much the same way in people with and without major depression, a disorder in which diagnosis depends on subjective reports. There's no reason we can't better understand mood at both a psychological and biological level, and I'm optimistic that this improved understanding will lead to some of the new treatments we so badly need.

SHAN SIDDIQI, BRIGHAM AND WOMEN'S HOSPITAL

As a neuropsychiatrist, I can explain some causes of my patient's memory disorder, and I can usually offer effective treatments for their mood disorder—but the converse is not satisfyingly true. How will we fill those gaps? First, we might need to recognize a fallacious dichotomy: either the disorders are

fundamentally similar, meaning both are brain disorders, or fundamentally different, meaning one is of the brain and the other of the mind. Both models are useful, but both are wrong (like all models). This ancient question plagued philosophers from Buddha to Descartes. But it was more relevant when epilepsy was attributed to evil spirits, all lung disorders were called "pleuritis," and the pineal gland was considered the seat of the soul. Medicine has since moved on, pragmatically adopting empiricism over Descartes' rationalism. It would be absurd to debate whether COVID-19 and lung cancer are fundamentally similar or different. Similarities and differences are both obvious. Perhaps more importantly, we have vastly different treatment approaches for COVID-19 and lung cancer, but also commonalities—we can reduce mortality globally by vaccinating against infection, raising nicotine taxes and improving health-care access. So are mood disorders caused by social factors or biological factors? The answer is "both." But the question is wrong. I won't live to see the full chain of causality for either mood or memory disorders. For now, we can make progress by mapping which parts of that chain can be targeted in which situations to reduce real suffering in real people.

ERIC TURKHEIMER, UNIVERSITY OF VIRGINIA

Enthusiasts and skeptics of a biological explanation of human behavioral traits face twin challenges. For us skeptics, the trick is to avoid ghost-in-the-machine dualism. No soul, no free will, no mind, unless it is unambiguously rooted in physical, evolved organic reality. Enthusiasts must avoid tautological materialism. I get it: all thinking, wanting and believing is done with the brain, and if all you have to say is that therefore aphasia and religious devotion are all equally "in your brain," fine, but it gets you nowhere. All assertions that a complex behavior is "biological" should be paired with an example of something that is not. In 1998, I imagined two silent people: a person who had a stroke in the left frontal lobe that resulted in a deep aphasia, and a monk who had taken a vow of silence as a matter of religious devotion. No dualism: Both silences are in their brains, somehow. Why do they seem so different? The answer is mostly about entities and the language we use to define and describe them. Aphasia, as a category, corresponds to a class of neurological events resulting from stroke. There is no class of neurological events that corresponds to devotional silence.

Problem

solving

We found a major flaw in a scientific reagent used in thousands of neuroscience experiments—and we're trying to fix it.

ILLUSTRATION BY DANIEL LIÉVANO

As part of that ambition, we launched a public-private partnership to systematically evaluate antibodies used to study neurological disease, and we plan to make all the data freely available.

BY MONA ALQAZZAZ, COMMUNITY ENGAGEMENT SCIENTIFIC FELLOW AND ALLIANCE MANAGER, YCHAROS AND THE STRUCTURAL GENOMICS CONSORTIUM; AND ALED EDWARDS, PROFESSOR OF MOLECULAR GENETICS AND MEDICAL BIOPHYSICS, UNIVERSITY OF TORONTO

The most common genetic cause of frontotemporal dementia is a mutation in the gene C9ORF72, first pinpointed in 2011. This amazing discovery was followed by years of confusion about the cell biology of the protein C9ORF72 encodes. Some papers reported that it was localized to the nucleus, others to a range of different organelles, still others to the cytoplasm. And without a basic understanding of where the protein acts, there was no way to move toward developing a treatment for frontotemporal dementia.

In 2019, working with a group of colleagues, we decided to look into the problem, focusing first on characterizing the antibody reagents other scientists had used to localize the protein. What we found was distressing. Using cells lacking C9ORF72 as controls, we discovered that not a single antibody reagent used in any of the published studies actually worked as advertised—they all bound to other proteins in addition to the target. In short, all the studies published using these C9ORF72 antibodies were potentially flawed. We also tested other commercial antibodies that had not, to our knowledge, been used in published research, and we found a few that were highly selective. These antibodies revealed that C9ORF72 is localized to the peri-lysosomal region and

mostly expressed in microglial cells, discoveries that have since been replicated.

The C9ORF72 example illustrates a widespread problem. Antibodies are one of the most commonly employed reagents in molecular research, used to identify single proteins in a cell's complex mixture. But scientists have known for decades that they can be inaccurate, binding to more than just the protein of interest. Publications that unknowingly use inaccurate antibodies can compound the issue, making it difficult to reproduce scientific results and raising questions about the validity of some preclinical drug studies.

Despite the seriousness of the problem, the field lacks a systematic way to characterize antibody accuracy. Quantifying how precisely an antibody highlights its target—its selectivity and specificity—is expensive and time consuming. The gold-standard approach is to compare cells expressing a target protein and those genetically modified to lack the target protein. Though many manufacturers do some knockout testing, the process is too expensive to apply to all antibody products. Most labs lack the requisite technologies, time or funding to rigorously characterize antibodies. As a result, most homemade or commercial antibodies are not subject to strict testing, a serious structural failure in the antibody ecosystem.

To address these structural issues, a team of us in academia and industry has launched an effort to systematically characterize widely used antibodies, employing knockout cells and tissues as controls whenever possible. Known as YCharOS, this public-private partnership has initially focused on antibodies used to study neurological conditions, including Alzheimer's disease, amyotrophic lateral sclerosis (ALS), autism, frontotemporal dementia and Parkinson's disease. We plan to eventually characterize high-performing antibodies for all human proteins.

In a pilot project, we tested 614 antibodies to each of 65 proteins linked to ALS, Alzheimer's disease and Parkinson's disease, using knockout cell lines we generated or collected from academic and commercial partners. The results were sobering. Many commercial antibodies did not perform as expected—60 percent of the antibodies were not specific to their intended target. Given this figure, we predict that as much as 20 percent of the figures in publications using these antibodies are in question. On the positive side, we estimate that well-performing antibodies are already available for more than 50 percent of protein targets. With this knowledge, antibody makers can focus their development efforts on the other 50 percent.

> "The open-science nature of the project was essential in making it work—companies and others are willing to participate because the data are open."

In just three years, YCharOS has already produced valuable information for the neuroscience community. But we could make faster progress with greater community contribution—partic-

ularly with community-supported sharing and generation of knockout cell lines.

All the data we generate will be publicly available after internal scientific review. Indeed, the open-science nature of the project was essential in making it work—companies and others are willing to participate because the data are open. Companies in the partnership have donated more than 800 antibodies for our pilot project, for example.

To date, YCharOS has published 78 antibody characterization reports for targets associated with ALS, Alzheimer's disease, autism, frontotemporal dementia and Parkinson's disease. These reports are available in the public domain at Zenodo, a general-purpose open-access data repository operated by the European Organization for Nuclear Research, CERN. To increase awareness, many of these reports have been shared via F1000 or more traditional publications, which enables them to be indexed on PubMed. Our protocols for immunoblotting, immunofluorescence and immunoprecipitation, developed in consultation with an expert group of academics and industry representatives, are also now publicly available. The protocol for immunohistochemistry is in development.

The reaction from commercial providers has been outstanding. When antibodies perform inadequately in the mutually agreed-upon standard operating procedures, their makers either remove the antibodies from the market (18 percent of antibodies tested) or alter their recommended usage (37 percent of the antibodies tested). We are also seeing an increase in the use of well-performing antibodies in the literature.

Manufacturers have also adjusted their research and development pipelines to meet the gaps that YCharOS has identified. For example, when it was clear that no existing antibodies were available for the proteins PRKN and SMOC1, two companies produced new well-performing antibodies. Similarly, based on feedback from YCharOS, many manufacturers are gearing up their pipelines to produce new recombinant antibodies to HTT, the protein implicated in Huntington's disease.

Looking forward, the field needs to address the availability of knockout cells, a key bottleneck. We could assess more antibodies if individual scientists shared their lines with YCharOS and if funders supported concerted community efforts to generate knockout lines. Scientists interested in a specific target could also include a budget line item to support the characterization of all commercial renewable antibodies for that target in each of many applications—we estimate this would cost roughly $50,000 per target. All these efforts would have a lasting impact because the characterization is focused on renewable and recombinant antibodies.

We have a clear scientific road map to identifying well-performing antibodies for all human proteins and an organizational framework to carry out this task. If the community works together, we will get there faster.

Women are systematically under-cited in neuroscience. New tools can change that.

PHOTOGRAPH BY RICHARD DRURY

An omitted citation in a high-profile paper led us to examine our own practices and to help others adopt tools that promote citation diversity.

BY ANNE CHURCHLAND, PROFESSOR OF NEUROBIOLOGY, UNIVERSITY OF CALIFORNIA, LOS ANGELES; AND FELICIA DAVATOLHAGH, POSTDOCTORAL RESEARCHER, UNIVERSITY OF CALIFORNIA, LOS ANGELES

One day last year, we opened a journal web page and were excited to find a long-awaited paper about a particular brain area's role in a cognitive computation. The authors, based on talks they had given at conferences, had a different take on this brain area compared with another high-profile paper, and we were curious to learn more.

But our hopes were soon dashed: The new paper didn't cite the previous one, published only a year before and with female first and last authors. The authors of both papers later discussed the omitted citation on Twitter/X, but the issue remained unresolved. It didn't seem that the authors of the later paper had omitted mention of the earlier one on purpose, so why was our female colleagues' highly relevant work left out?

Research suggests the answer may lie with systematic citation bias—a trend that persists even as the number of women in academia, including in neuroscience, continues to grow. The proportion of papers authored by women (first or last author) across five broad-scope neuroscience journals increased from 36 percent in 1995 to 50 percent in 2018. Yet first and last male authors are increasingly over-cited, according to a study published in Nature Neuroscience

"It is urgent that we address this issue, given that citation metrics are often used as a measure of a researcher's impact and productivity in their field and can influence their hiring, promotion and invited-speaker opportunities."

in 2020. And the gap between observed and expected gender proportions in neuroscience citations is widening at a rate of 0.41 percentage points per year, even after accounting for year of publication, number of authors, article type and first/last author seniority.

The imbalance is largely driven by gendered citation practices—male authors are more likely to over-cite male authors and less likely to cite women-led work. A subsequent analysis that focused on reference lists within cognitive neuroscience found a similar pattern of bias, as did a study examining the intersection of race, ethnicity and gender.

It is urgent that we address this issue, given that citation metrics are often used as a measure of a researcher's impact and productivity in their field and can influence their hiring, promotion and invited-speaker opportunities. Fortunately, new tools—such as cleanBib, anneslist and others—exist to help researchers counter citation bias across genders and other demographic divides.

These tools have some limitations, such as a lower accuracy for transgender, nonbinary or intersex authors, but they are of great value to our field. Actively performing literature searches to improve citation balance can improve a manuscript's scholarliness, as it may lead authors to identify work that was overlooked initially.

CleanBib analyzes any reference list and quantifies the proportion of men and women authors, as well as the proportion of authors from other underrepresented groups, such as Black and Latino neuroscientists. It then generates a citation diversity statement that breaks the bibliography down by gender and ethnicity and can be placed before a paper's references section. A 2021 analysis found that cleanBib has already had an impact. Papers from scientists that have used the tool cite a larger proportion of women than the average rates in five top neuroscience jour-

nals: *Nature Neuroscience, Neuron, Brain, Journal of Neuroscience* and *Neuroimage.*

Additional tools—such as anneslist, a list of women scientists in different subfields of neuroscience; Cite Black Authors, a database of research by Black academics; and Diversify STEM Conferences, a list of researchers from underrepresented groups across fields—can also help authors identify relevant researchers and improve citation balance.

Asking scientists to include a citation diversity statement is one way to promote an awareness of citation bias and to encourage them to use these tools and acknowledge the scientific contributions of women and people of color. Cell Press was among the first publishers to invite authors to create an inclusion and diversity statement, and it found that more than 40 percent of authors opted to do so. This initiative was paused as of late 2023, but we hope it eventually continues.

Including citation diversity statements in preprints would provide even more opportunity for authors to receive feedback on their bibliographies. Ideally, these practices could be implemented along the lines of data-sharing and open-science policies. Researchers who, like us, have used cleanBib and other tools could share their experience and expertise.

Including a diversity statement in research publications is an actionable goal for everyone, from graduate students writing their first paper to principal investigators reviewing manuscripts from their lab. More broadly, investigators reviewing papers for journals should pay attention to whose work is cited and offer suggestions if work is excluded. We all must play a role in determining whose work is acknowledged and valued while moving the field toward reference lists that accurately represent the increasing diversity of the neuroscience field.

We were among the first researchers to add diversity statements to papers, and recently we ran a hands-on workshop at the University of California, Los Angeles on citation bias, demonstrating how to use cleanBib to quantify bias in gender, race and ethnicity in reference lists and exploring its strengths and limitations. Efforts like this are part of our longstanding commitment to leveling the playing field in science and making it more welcoming to newcomers of all kinds. Looking back on the female-female paper that was omitted in the recent manuscript, we can't help but wonder whether these tools could have avoided the unintentional oversight.

At the credit crossroads: Modern neuroscience needs a cultural shift to adopt new authorship practices

ILLUSTRATION BY VAHRAM MURADYAN

Old heuristics to acknowledge contributors—calling out first and last authors, with everyone else in between—don't work well for large collaborative and interdisciplinary projects, yet they remain the default.

When it comes to how neuroscientists assign credit in academic writing, the field is at a crossroads. Neuroscience has roots in biology and psychology, which have traditionally favored smaller collaborations, and it relies on simple heuristics, such as authorship order, to assign credit: The first author did all the work, the last author supervised, and a few folks in between played various (smaller) supporting roles. But as neuroscience broadens to embrace cognitive and computational neuroscience, artificial intelligence, big data and more, the field is venturing into a Wild West of large consortium science.

With larger, collaborative and increasingly interdisciplinary efforts, the question of who really gets credit for a given scientific output becomes much more complex—and established cultural norms no longer work. Research contributions in neuroscience and psychology are more numerous, more varied and more specialized than ever, and the increasing adoption of open-science practices calls for even more nuanced credit assignment. How can we push forward into this brave new "big science" world while properly recognizing these contributions, both practically and socially?

Researchers have developed several potential solutions to this problem, including new ways to assign both the over-

BY MEGAN PETERS, ASSOCIATE PROFESSOR OF COGNITIVE SCIENCES, UNIVERSITY OF CALIFORNIA, IRVINE

all amount of credit and the type of credit each contributor should receive. I've suggested a few ideas below. But so far, these clever guidelines have failed to replace the first-author, last-author convention. Why? Personally, I think a true solution will require more than developing ontologically satisfying systems for credit assignment and integrating them into our publishing systems. Instead, we must find better ways to integrate these systems with established cultural norms and—importantly—our very human need for simple, interpretable heuristics.

To see why, let's lay out the scope of the problem.

First, what is the appropriate amount of credit to assign to a contributor? Most scientists recognize that assigning a monolithic amount of credit based on the first-author, last-author convention often fails to reflect the volume or criticality of work done by others on the list. We've tried a few fixes in neuroscience, including the increasing prevalence of co-first and co-last author arrangements. But I think we all know in our hearts that this doesn't address the core challenge: People still fight over who is really first on co-first-author lists. (Some citation software still lists "Smith et al., 2005" in American Psychological Association [APA]-style in-text citations, even if the full citation is Smith*, Jones*, Lee, & Lau, with the asterisks denoting equal contribution.)

And how do we get the credits right for large, multi-lab collaborations, in which there are several trainees sharing the work equally and several equal co-principal investigators?

We could take inspiration from other fields, including particle physics and public health, that have applied a remedy at one extreme end of a spectrum: treating an "author" as essentially everybody in a consortium who ever touched some large piece of equipment or who contributed to a particular dataset. As a result, some papers have hundreds or even thousands of authors (see, for example the Large Hadron Collider paper with 5,154 authors, or the COVID-19 vaccination paper with 15,025 co-authors), or circumvent first-authorship stardom by putting the name of the consortium as first in the author list.

Some neuroscience consortia have adopted similar approaches, such as the International Brain Laboratory, which offers complex, community-developed rules for authorship. But some major concerns prevent the blanket adoption of these practices, especially in groups of tens rather than hundreds or thousands of authors. Chiefly, large-list approaches can backfire in some fields, research shows: Rather than spreading the wealth as intended, they may obscure credit, meaning no one gets the recognition they deserve.

In 2022, for example, Clarivate, a large player in data analytics that ranks "highly cited researchers," began excluding papers with more than 30 authors from its calculation. Because the median number of authors per paper varies highly from field to field, this issue disproportionately affects fields with larger author lists, including neuroscience: As of 2018, roughly 80 percent of papers in psychiatry and psychology—and 91 percent and 74 percent, respectively, in computer science and physics—had one to five authors, compared with about 50 percent of papers in biology and neuroscience. Because modern neuroscience is heavily

"The challenge we're facing in credit assignment isn't just ontological—it's deeply cultural, psychological and practical."

influenced by fields with fewer authors, we might find it particularly difficult to grapple with longer author lists.

Some have proposed new indices to augment existing practices by indicating "how much" a given author contributed, but cultural uptake of such metrics in our field has unfortunately been slow to nonexistent. A second, perhaps more challenging problem is how to recognize different kinds of contribution—collecting data, performing analyses, conceptualizing a project and acquiring funding—that are qualitatively different from one another.

One solution is the CRediT (Contributor Roles Taxonomy) system, introduced in 2015 as an attempt to provide transparency and accountability in authorship type attribution. This is the system that asks, when you submit to a journal, to tick up to 14 boxes saying who did what: conceptualization, methodology, software, validation, funding acquisition, writing and so on, with the results typically displayed in an "Author contributions" section of the manuscript. Championed as a way to diminish the use of authorship order in assigning credit, the CRediT system has been adopted by more than 40 publishers and journal families, including Public Library of Science (PLOS), Cell Press and Elsevier, and was recently adopted as an American National Standards Institute/National Information Standards Organization standard.

Unfortunately, though, I don't think CRediT works as well in practice as we would hope. First, CRediT's limited scope and narrow focus on predetermined traditional authorship roles may mischaracterize the diverse and ever-evolving range of contributions in larger, consortium style research—especially as we move to more complex (and laborious) data annotation and sharing. (For more, see Maryann Martone's piece from The Transmitter's "Open neuroscience and data-sharing" essay series.) The restriction of applying a CRediT-type system only to authors also obscures other important contributions that may not surpass a journal's threshold for authorship, such as from technical staff. The one-size-fits-all categories also don't always work well for opinion, educational materials, or perspective style pieces—the kinds of pieces we have typically written in large, multi-author groups at Neuromatch, for example.

What about alternatives? Many smart, driven meta-scientists have developed promising innovations beyond CRediT: the expanded Contributor Role Ontology; authorship matrices; storyboarding approaches; and the Contributor Attribution Model, "an ontology-based specification for representing information about contributions made to research-related artifacts." But uptake of these tools—and a corresponding shift in how we collectively think about and use credit assignment in practice—has proven even more achingly slow than the meaningful adoption of CRediT in daily life. This is also especially important for trainees, whose careers might be most negatively impacted by attempts to adopt idealistic but culturally challenging quick fixes, such as multiple co-first authors or "Consortium et al." approaches.

Here's a perhaps less obvious but, in my opinion, crucial problem: As the most widely adopted system, CRediT formats author contributions to be machine readable and accessible via API, in addition to being included as "Author contributions" statements on individual papers. This formatting makes CRediT-based contributions useful for meta-analyses, for example, but less useful in aggregate for any end-user scientist like you or me. Sure, we can read all those author contribution statements on every PDF, I suppose. But it is difficult—whether we are trying to evaluate colleagues or applicants in hiring or promotion decisions or simply curious—to use CRediT-based data to figure out what kinds of contributions a particular person has made across their entire career, or even just recently.

ORCID profiles can technically integrate CRediT information from ORCID-linked papers, but they don't summarize this information in a useful way. (Few people check someone's ORCID profile page anyway, and even the account owner doesn't have easy access to their own aggregate metrics.) So rather than relying on these contributions, it's simply much easier for us all to continue characterizing our fellow scientists by counting their first-authorships and high-profile outlets on a CV or looking at their h-index on Google Scholar, even though we all know better.

To put it plainly, the challenge we're facing in credit assignment isn't just ontological—it's deeply cultural, psychological and practical. And so far, simply recognizing that the issue exists hasn't been enough to make us change our ways.

I think our path forward is clear: We need better methods for meaningfully integrating new credit assignment systems into our existing workflows in ways that make the information obvious, transparent and accessible in daily life. Heuristics become entrenched for a reason—we keep falling back on the "first author, last author" shortcut because it's easy and makes sense—so let's make ourselves some cheat sheets that better reflect our values. ORCID profile contribution badges across all linked papers could be a good start; a new category of summary statistics to accompany the h-index and i-index section on Google Scholar or cute graphical ways to display our identities as researchers on our CVs might also help. I'm sure you have other ideas, too. But no matter which version actually catches on, the collective goal should be a shift in our focus toward prioritizing the sociological utility of credit assignment, rather than simply capturing the data. Hopefully, a cultural shift will come along for the ride, and we'll all be better at appropriately recognizing the valuable diversity of "big science" authorship contributions in modern neuroscience.

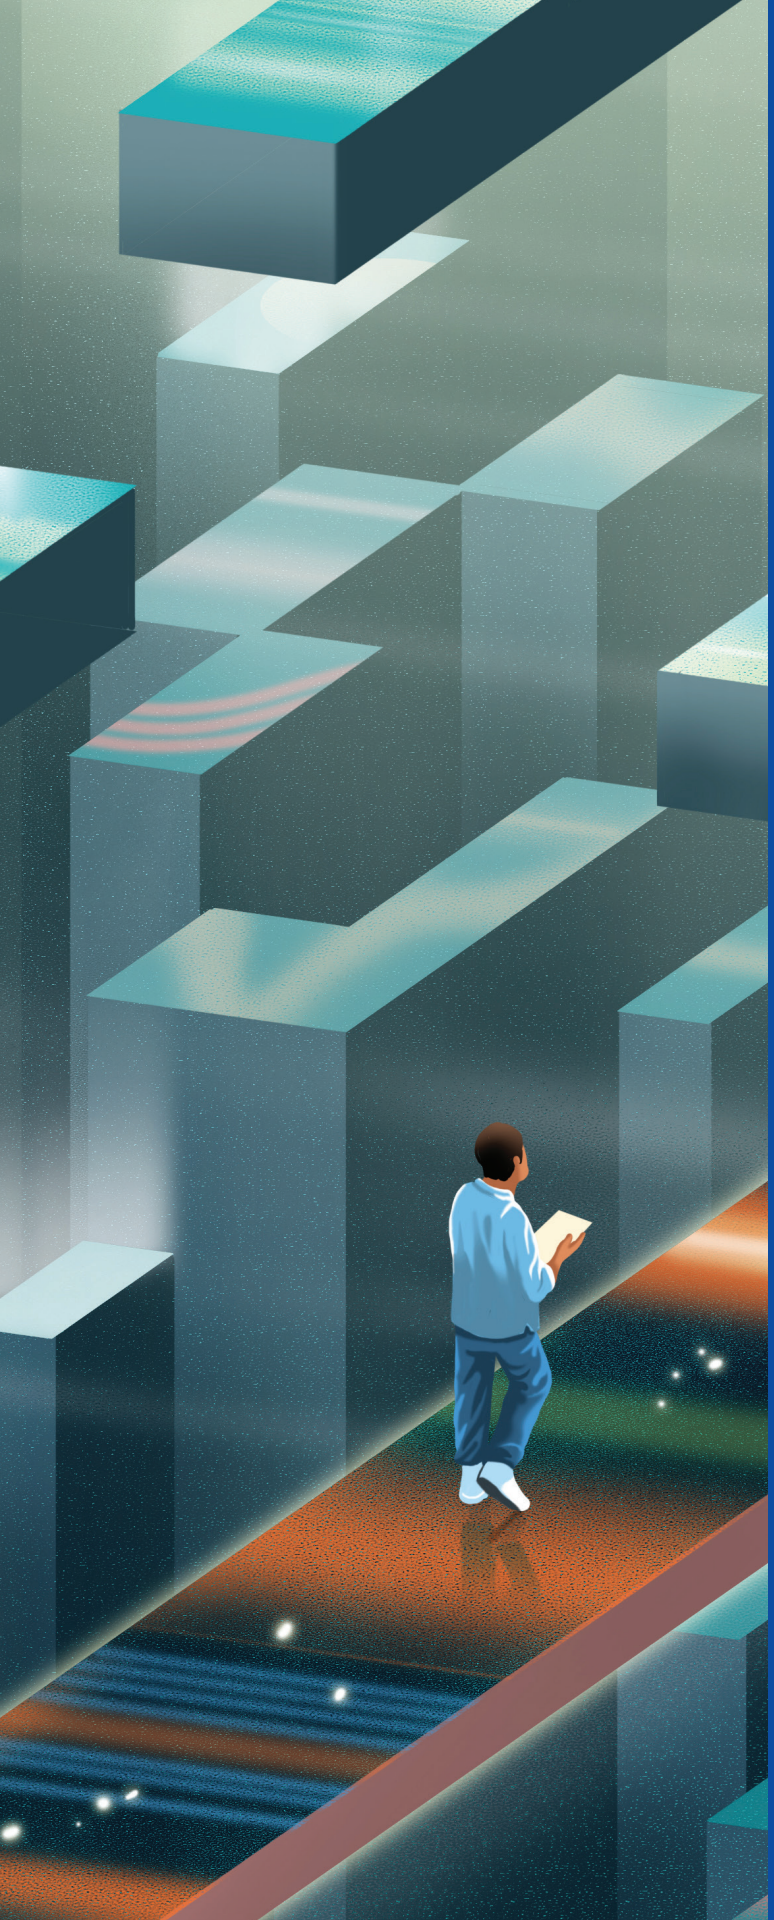

How long-read sequencing will transform neuroscience

thetransmitter.org/whole-genome-sequencing/how-long-read-sequencing-will-transform-neuroscience/

New technology that delivers much more than a simple DNA sequence could have a major impact on brain research, enabling researchers to study transcript diversity, imprinting and more.

BY TYCHELE TURNER, ASSISTANT PROFESSOR OF GENETICS, WASHINGTON UNIVERSITY SCHOOL OF MEDICINE

Defining

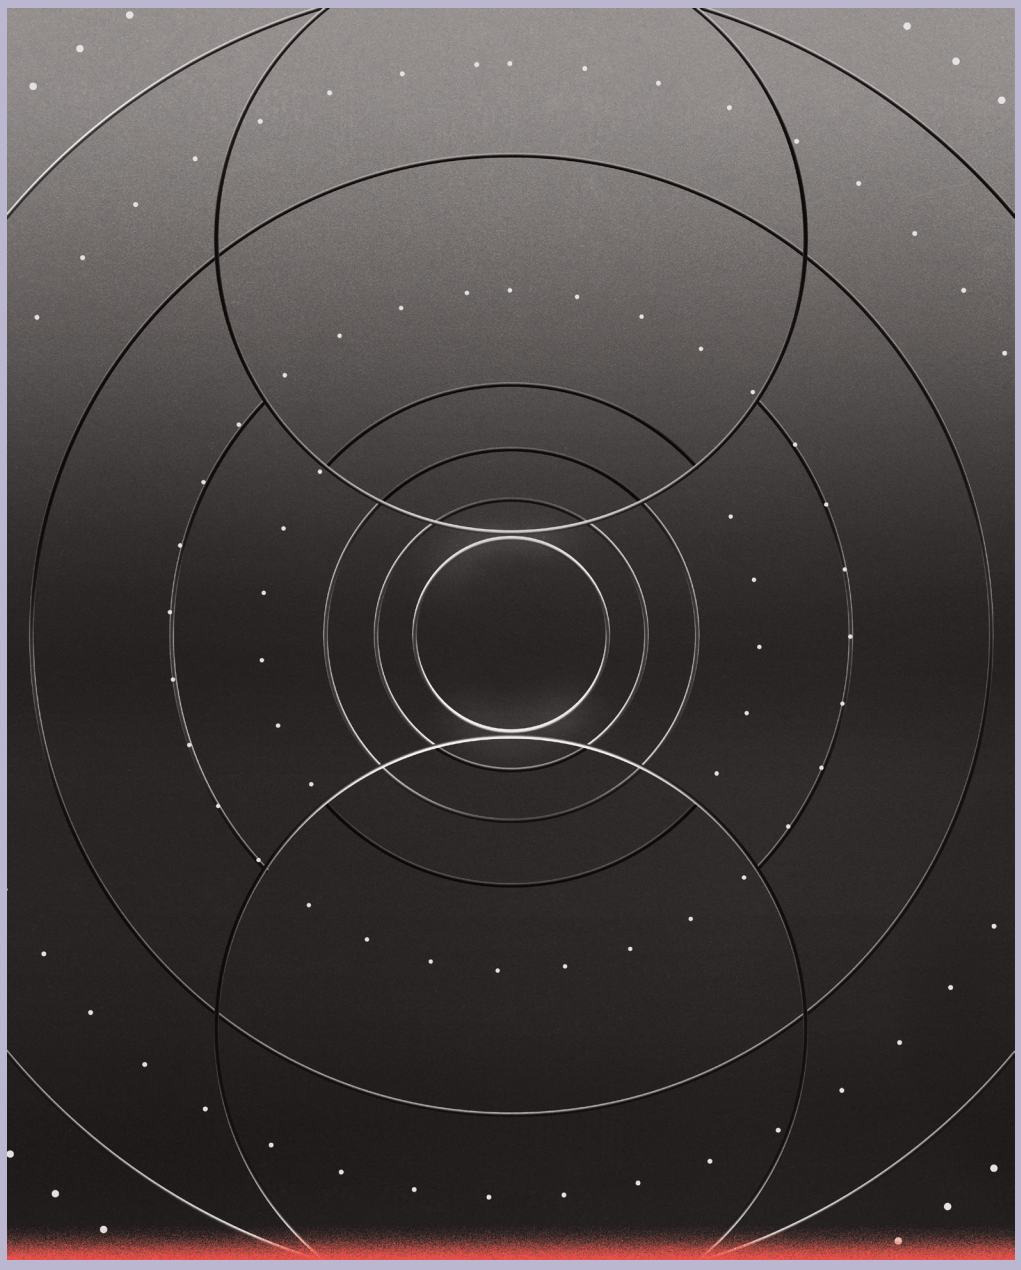

representations

This series explores the often-fuzzy concept of representation and the different ways researchers employ the term.

What are we talking about? Clarifying the fuzzy concept of representation in neuroscience and beyond

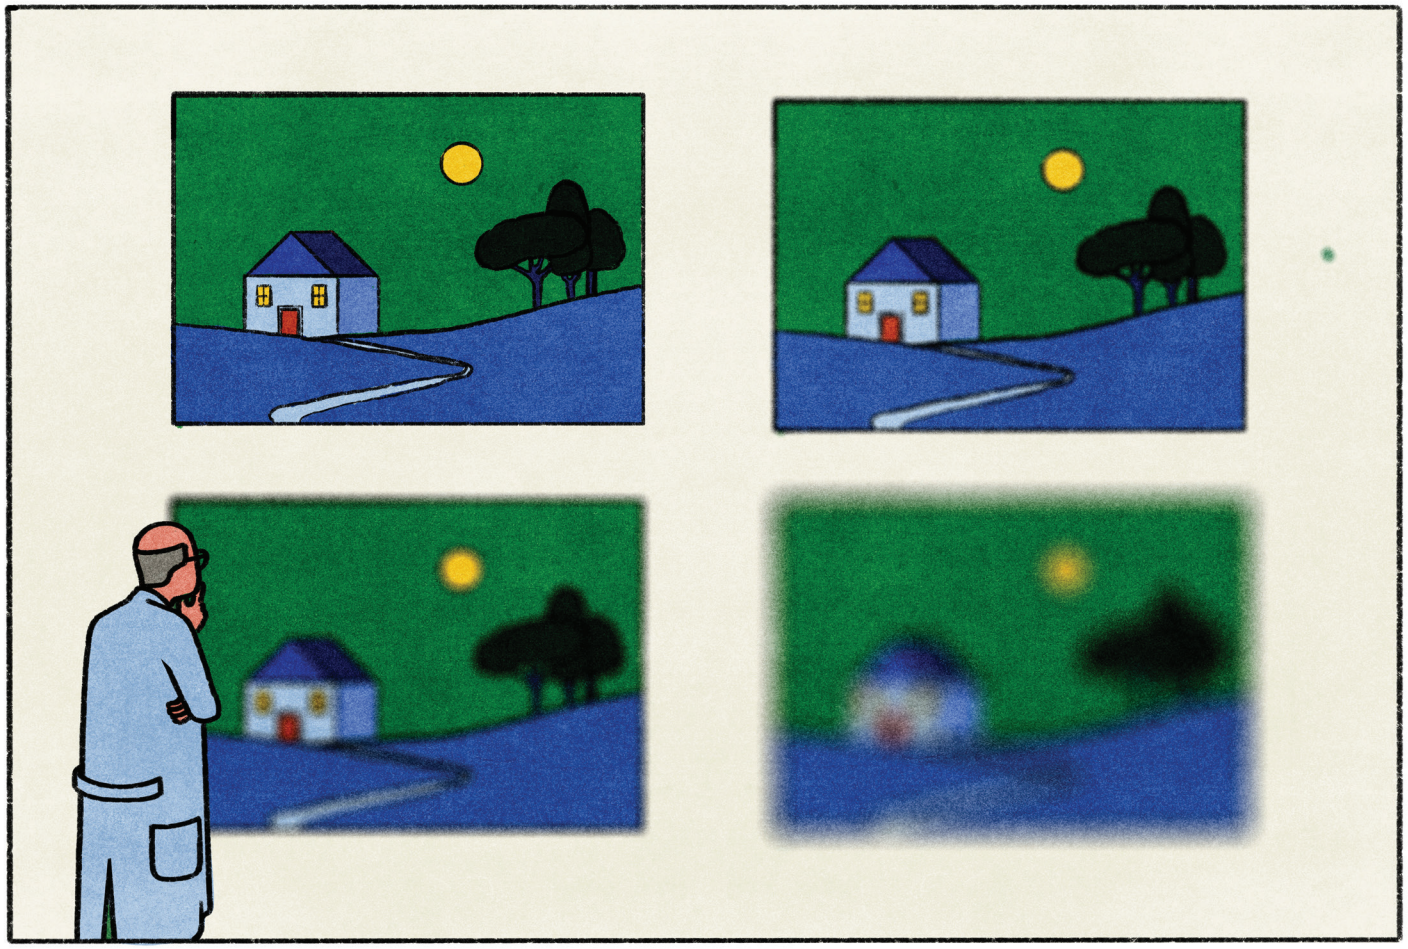

ILLUSTRATION BY KLAUS KREMMERZ

To foster discourse, scientists need to account for all the different ways they use the term "representation."

The notion of representation in brain and cognitive sciences is ubiquitous, vitally important, and yet fuzzy. This holds both within neuroscience and beyond, including in cognitive science, artificial intelligence, linguistics, psychology and philosophy of the mind. What people mean when they use the term varies considerably, ranging from a simple correlation between a neural response and a stimulus to a true offline model of the world. Indeed, the members of our group hold many different, often opposing, views on how best to define the concept of representation. To enable clearer usage and facilitate discussion, we hope as a group to develop a catalog of the various uses.

For many, representation is central to the very idea of a science of the mind. This view was already well enough established in 1983 for linguist Noam Chomsky to write: "It is fair to define cognitive psychology as the study of mental representations—their nature, their origins, their systematic structures, and their role in human action." Consider also this entry from the Encyclopedia of Philosophy: "Mental representations are the coin of contemporary cognitive psychology, which proposes to explain the etiology of subjects' behavior in terms of the possession and use of such representations." A common application is within compu-

BY FRANCIS T. FALLON, ASSOCIATE PROFESSOR OF PHILOSOPHY, ST. JOHN'S UNIVERSITY; TOMÁS J. RYAN, ASSOCIATE PROFESSOR OF NEUROSCIENCE, TRINITY COLLEGE DUBLIN; JOHN W. KRAKAUER, PROFESSOR OF NEUROLOGY, JOHNS HOPKINS UNIVERSITY AND THE RPPF GROUP

tational models of cognition, where models can independently draw on discrete representations, recombine them in a compositional manner and operate on them to enable planning, reasoning and problem-solving.

Cognitive scientists use the term representation to posit an entity that lies at the basis of cognition, but neuroscientists instead use it to indicate that behaviorally relevant information is detectable in single neurons, circuits and neural populations. This more casual or correlational usage of the term makes fewer theoretical commitments, but as a result it is often not clear what additional work it does beyond referring again to the mere association itself.

These types of correlations are quite easy to detect, and new findings question whether all of them should be referred to as representations. For example, recent studies have found that activity across the brain, even in early sensory neurons or areas, correlates with a variety of properties, including ongoing actions, choices and behavioral engagement. Is the activity of these neurons causally relevant to cognitive function in every case—meaning, for example, does that neural activity actually contribute to the animal's decision—or is it just epiphenomenal? Regardless, this recent work showing correlations everywhere raises questions about whether simple correlation is a strong enough basis for representation.

Even in the realm of cognition, there have been attempts that seem to forego mental representations altogether and instead emphasize connectionist and dynamical systems views. Such views often, but not always, eschew the requirement for an overt representation as a vehicle with content that exists in some correspondence with the world. Instead, behavior is the output of a distributed network operating on an input. A central pattern generator, for example, can drive locomotion without having to explicitly represent a leg anywhere inside its circuitry. Larger networks could operate as a scaled-up version of this.

> "We propose that a catalog of the different senses of representation would greatly facilitate communication in cognitive science."

Proponents of such views may still want to apply the term "representation." For example, single cells in a central pattern generator that correlate with leg kinematic variables such as position and speed could be said to represent them. But mere correlation does not seem like a compelling basis for attributing representation. There may also be a parallel lesson to draw in AI: Peering into artificial neural networks and finding responses that look like the ones found in the primate cortex has led some to posit similar representations. But do similar responses in a specific layer mean that a similar overall representation is present? Again, as in the single-cell case, it is not clear what applying the term "representation" would mean, beyond indicating that task-relevant information can be found distributed throughout the network.

What to do? One temptation is to remain vague about what the term "representation" means, to avoid impeding its use in any discipline-specific research. Another temptation is to deflate its meaning to its most minimal form: Use it whenever there is information in neural data, single neurons or neural populations that correlates with task features in the broadest sense. Alternatively, perhaps the term should be reserved for only the most maximal form: referring to a discrete type of true neural state with flexible abstract content, which can be activated absent its original cause—i.e., a representation that is in fact used as a representation.

Some argue that this last type of representation should be considered categorically distinct from other, looser senses of the term. It might even be that this is what distinguishes uniquely human thinking. Still others, largely in allergic response to this more fleshed-out notion of representation, have called for the term's elimination from the scientific literature (A 1990 paper by Walter Freeman and Christine Skarda is an unambiguous example, but see also 2019 work by Rafael Nunez and his colleagues, a 2019 article by Romain Brette, a 2017 paper by Daniel Hutto and Erik Myin, Alva Nöe's 2006 book on perception and Anthony Chemero's 2011 book on cognitive science).

Because the options are either problematic or controversial, we propose that a catalog of the different senses of representation would greatly facilitate communication in cognitive science. Such a taxonomy would enable scientists to choose descriptors of varying levels of specificity and inspire researchers to more carefully consider when, how and whether they use it, and to communicate what they mean more explicitly. The taxonomy need not favor any particular theory of representation, nor even assume its existence; rather, it would help set the terms for discussion surrounding an otherwise ambiguous and confusing term. We submit that there is a growing will in neuroscience, philosophy and the cognitive sciences to engage in just such a project (see also a 1987 essay by Ernst von Glaserfeld and a June 2023 paper by Luis Favela and Edouard Machery).

The RPPF is hosted at Trinity College Dublin in Ireland and is funded by an Institutional Strategy Support Fund grant from the Wellcome Trust.

When do neural representations give rise to mental representations?

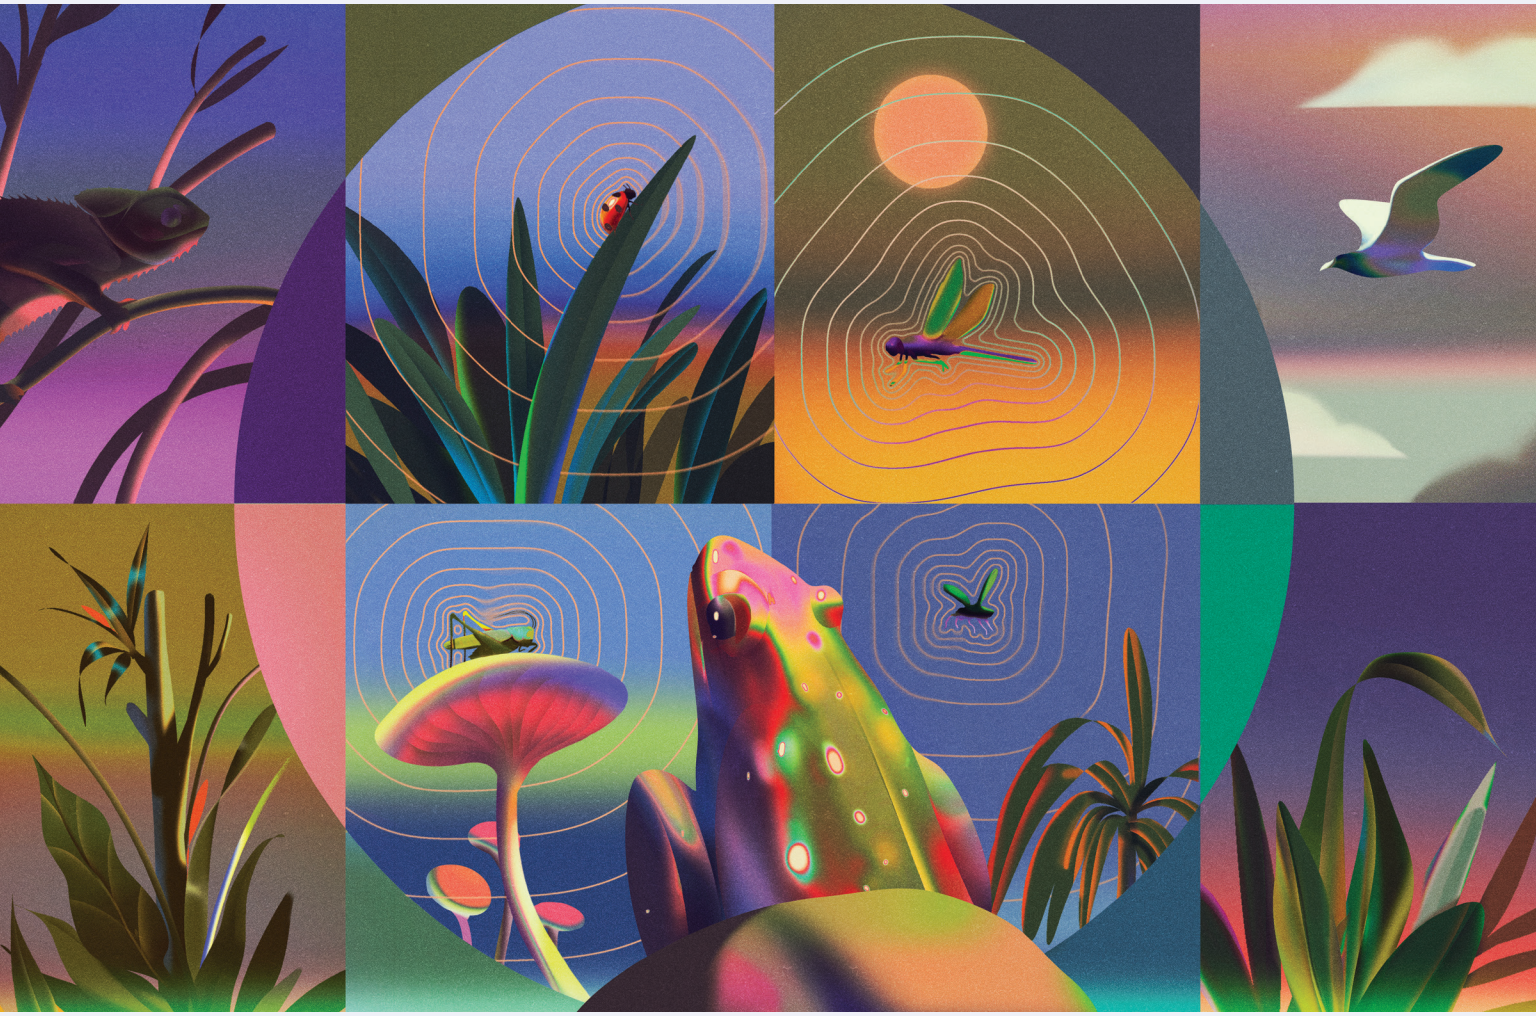

ILLUSTRATION BY IBRAHIM RAYINTAKATH

To answer this question, consider the animal's umwelt, or what it needs to know about the world.

It is often said that "the mind is what the brain does." Modern neuroscience has indeed shown us that mental goings-on rely on and are in some sense *entailed by* neural goings-on. But the truth is that we have a poor handle on the nature of that relationship. One way to bridge that divide is to try to define the relationship between neural and mental representations.

The basic premise of neuroscience is that patterns of neural activity carry some information—they are *about something*. But not all such patterns need be thought of as representations; many of them are just signals. Simple circuits such as the muscle stretch reflex or the eye-blink reflex, for example, are configured to respond to stimuli such as the lengthening of a muscle or a sudden bright light. But they don't need to internally *represent* this information—or make that information available to other parts of the nervous system. They just need to respond to it.

More complex information processing, by contrast, such as in our image-forming visual system, requires internal neural representation. By integrating signals from multiple photoreceptors, retinal ganglion cells carry information about patterns of light in the visual stimulus—particularly edges where the illumination changes from light to dark.

BY KEVIN MITCHELL,
ASSOCIATE PROFESSOR OF
GENETICS AND NEUROSCIENCE,
TRINITY COLLEGE DUBLIN

This information is then made available to the thalamus and the cortical hierarchy, where additional processing goes on to extract higher- and higher-order features of the entire visual scene.

Scientists have elucidated the logic of these hierarchical systems by studying the types of stimuli to which neurons are most sensitively tuned, known as "receptive fields." If some neuron in an early cortical area responds selectively to, say, a vertical line in a certain part of the visual field, the inference is that when such a neuron is active, that is the information *that it is representing*. In this case, it is making that information available to the next level of the visual system—itself just a subsystem of the brain.

Usually, patterns of neural activity across local populations represent such information. And crucially, the populations and circuits that *interpret* such representations are causally sensitive to the meaning of those macroscale patterns, rather than to the details of the instantaneous neural instantiations.

But does that activity of *representing* imply that the pattern is *a representation of* the type envisaged in cognitive science—i.e., a distinct cognitive object? At what point do intermediate neural representations give rise to high-level mental representations? When does distributed, devolved signal processing become centralized cognition? A good way to approach this question is to ask: When does it need to?

"When does distributed, devolved signal processing become centralized cognition? A good way to approach this question is to ask: When does it need to?"

The point of perceptual information processing systems is ultimately to allow the organism to know what is out in the world and what the organism should do about it. Any given scenario will involve myriad factors and relationships, including dynamically changing arrays of threats and opportunities—too many to be hard-coded into reflexive responses. To guide flexible behavior, sensory information and information about internal states have to be submitted to a central cognitive system so this information can be adjudicated over in a common space. Such a system can and should only operate with some kinds of information—information *at the right level.*

We don't want to think about the photons hitting our retina, even though our neural systems are detecting and processing that information. That's not the right level of information to link to stored knowledge of the world or to combine with data from other sensory modalities. What we need to think about are *the objects in the world* that these photons are bouncing off of before they reach our eyeballs. We, as behaving organisms, need to cognitively work with our perceptual inferences, not our sensory data.

For any organism, we can ask two related questions: What kinds of things *can* it think about? And what kinds of things *should* it think about? Here, an ecological approach is valuable.

Jakob von Uexküll introduced the concept of the "umwelt" of an organism, which might be translated as its "experienced world." This idea begins with the specific sensorium of the organism, which will vary among species and sometimes even individuals. Different organisms can smell different chemicals, for example, or detect different wavelengths of light or frequencies of sound, while being oblivious to many other factors in their environment.

But the umwelt also crucially entails valence, salience and relevance— it is a self-centered map of things in the environment that the organism can detect and that *it cares about.* Perception in this view is not the neutral, passive acquisition of information. It is an active process of sense-making, which results in a highly filtered, value-laden, action-oriented landscape of affordances. If an organism is thinking at all, these are the things that it needs to think about.

Using the word "think" in this context invites all kinds of anthropomorphism, of course, and risks inflating to the level of cognition what might be better framed in cybernetic terms. But if cognition can be defined as using information to adaptively guide behavior in novel circumstances, then perhaps we can think of a continuum from simpler control systems to the more abstract cogitation that humans engage in. That is, after all, the trajectory that evolution had to follow. And if we allow that many animals have some kind of central arena where the highest-level inferences become the objects of cognition, we can ask what such inferences could be about.

That question may be difficult or even impossible to answer for animals incapable of self-report. But we may be able at least to say what various organisms *can't* be thinking about. A nematode can't be thinking about objects far away from it because it has no means to detect them. Its cognitive umwelt is consequently limited to the here and now. A lamprey can't be thinking *about* types of objects because it doesn't have enough levels of processing to abstract the requisite categorical relationships. And a human baby can't think about next week because its cognitive horizon doesn't extend that far.

Each animal's potential cognitive umwelt—what it *could* think about—is thus limited not just by its sensory capabilities, but also by the levels of internal processing it has, as well as its capacity for long-term memory and long-term planning. But there are also active limits on what any organism *does* think about. All the low-level information processing gives rise to the objects of cognition, but in a highly selective, filtered fashion. High-level cognition is useful precisely because it ignores so much low-level detail—because of what's *not* on your mind.

Thus, only a subset of neural representations—the meaningful elements that are processed by various neural subsystems—rise to the level of *mental representations*—the elements of cognition. Moreover, only a subset of those mental objects—a varying subset, depending on circumstances—may be things that we need to think about *consciously*. It's thus too vague and all-inclusive to simply say, "The mind is what the brain does." Our mental goings-on are more likely entailed by a dynamically shifting, adaptively filtered subset of neural goings-on.

How-tos

HOW TO TEACH THIS PAPER:
This column by Ashley Juavinett guides educators and self-learners through recent seminal neuroscience papers.

FROM BENCH TO BOT:
In the "From bench to bot" series, neuroscientist and science writer Tim Requarth explores the promises and pitfalls of artificial-intelligence tools in writing.

How to teach this paper: ‘Neural population dynamics during reaching,’ by Churchland & Cunningham *et al.* (2012)

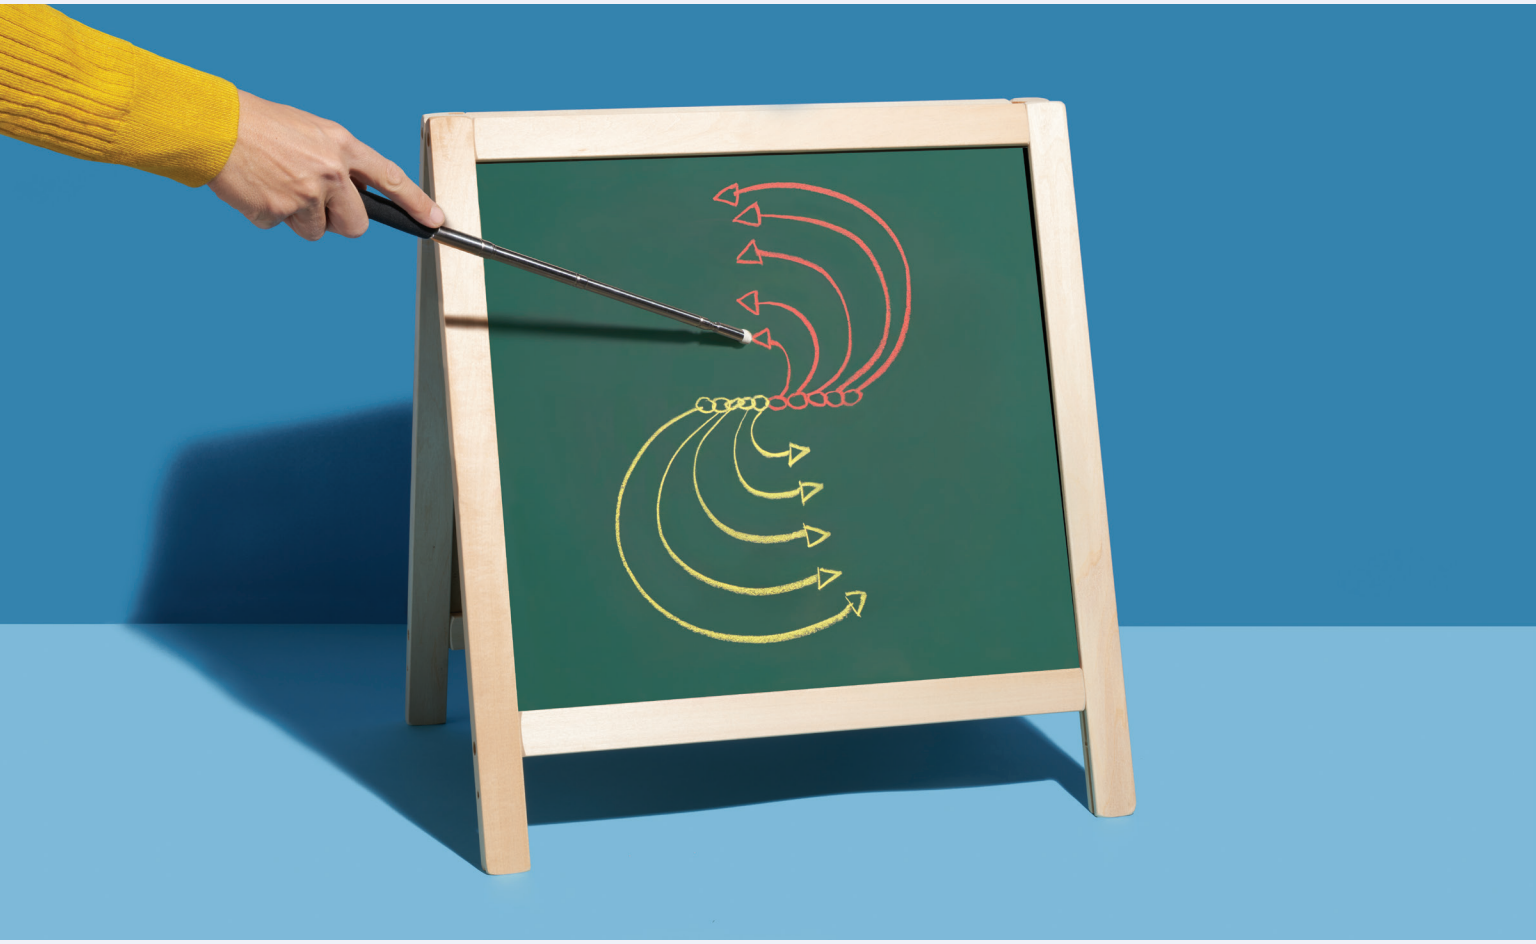

ILLUSTRATION BY MARGEAUX WALTER

This foundational paper, with more than 1,500 citations, is an important departure from early neuroscience research. Don't be afraid of the math in the first paragraph.

With almost 1,500 citations, "Neural population dynamics during reaching," published in *Nature* in 2012, is a foundational paper in systems and computational neuroscience. It presents the first evidence as well as a thorough theoretical footing for the idea that populations of neurons generate movement in a rather stunning way: with dynamics that rotate. This paper would be a great fit for systems or computational neuroscience classes, or an interesting applied example for a mathematics class. Given its focus on movement, it would also be a great choice for a class focused on motor systems or disorders of movement and could be tied to a discussion about brain-computer interfaces.

This paper is an important departure from early neuroscience research, which focused on individual neurons that represented information by changing their firing rate. Such neurons—largely studied in the visual system—were tuned to specific features, such as the orientation of a black and white bar, a direction of movement, or even pictures of Jennifer Aniston. But for a long time, movement researchers struggled to find neurons that reliably responded to only one feature of a movement. In the visual cortex, each neuron had a role; in the motor cortex, each neuron seemed to be doing everyone else's job.

BY ASHLEY JUAVINETT,
ASSOCIATE TEACHING PROFESSOR
OF NEUROBIOLOGY, UNIVERSITY
OF CALIFORNIA, SAN DIEGO

That's where Krishna Shenoy's team came in. Shenoy turned to more complex computational methods to help translate the neural activity in the motor cortex. Specifically, he championed a class of "dynamic" methods that ultimately became central to understanding how the brain produces movements and are the heart of "Neural population dynamics during reaching."

PAPER OVERVIEW

A great pedagogical feature of this paper is that it first takes a broad look at the neural activity underlying different movements as a means of pointing readers to a surprising observation: Neural population activity looks similar for different kinds of movement. Some movements, such as crawling, are visibly rhythmic, so this is a nice entry point to understand why movements may be generated by oscillating neural activity. By focusing first on the leech example given in the paper (perhaps bolstered by William Kristan's papers cited in that section), students can develop an intuition for what "rhythmic" means in this context.

From there, we can make the extension to oscillating firing rates in different kinds of movements, such as walking or reaching. Although these movements aren't oscillatory, their neural population response is—weird. Mark Churchland, John Cunningham and their colleagues use this observation to motivate a deep dive into more data. Shenoy gave an elegant, convincing talk about this transition from representational frameworks to dynamic systems in 2013, which I strongly recommend watching for background.

As an instructor, you might choose to stop at the first figure; it is rich enough in theory and analysis, yet very simply diagrammed (and with a cute leech drawing). But figures 2 through 5 show raw electrophysiology (single- and multi-unit arrays as well as electromyography) data from monkeys as they are reaching, which can be useful for students to see. They also illustrate another unexpected feature of the data: Rotations don't seem to depend on the actual movement direction and don't relate to the actual reach path of a monkey.

As the finale, the paper presents a model that illustrates that activity created by a generator model (which emulates a dynamical system) is more closely matched to the neural data than are EMG or two other kinds of simulated data models, which only take into account the features of movement, such as velocity, acceleration or direction. There is a lot packed into the final figure; as an instructor, how much you explain these models depends on your own comfort level and how relevant they are to your course. For most neuroscience courses, it's sufficient to say that when you simulate data, you're doing so because it allows you to generate activity with known properties or that can perform a particular task, so that you can see what happens when you pass such data through your analysis. For the Churchland paper, the only data that produce rotational dynamics are the neural data that the researchers have recorded and the generator model. In later work, David Sussillo, Churchland, Matthew Kaufman and Shenoy showed that a recurrent neural network would also spontaneously adopt brain-like, quasi-oscillatory patterns.

"Educators and students alike should not fear the math in this paper. Those equations are actually saying something very simple."

STICKY POINTS

Oh hello, math.

This paper does something that is either a very bold move for a neuroscience article or a very big faux pas, depending on your relationship with math: It includes equations in the first paragraph. From a student perspective, this can be quite daunting and an immediate reason to think, "There's no way I'm going to understand this paper." (Most biology students fear math.) But educators and students alike should not fear the math in this paper. Those equations are actually saying something very simple.

Let's start with Equation 1, which serves as the null hypothesis of the paper. It represents the prevailing idea in the field: A single neuron represents different features of a stimulus. In this equation, a neuron's firing rate is predicted by some combination of the parameters of a stimulus. For a limb movement, this might be the intensity of muscle contraction or the direction of the limb. As an instructor, it is up to you whether you would like to explain the nuances of this equation, and it may be a nice opportunity to review how mathematicians speak.

The second equation is the dynamical system equation, and even though it's shorter, it's conceptually more complex. Now we're trying to compute the firing rate of all neurons (the population code). The derivative of the population code is determined by some unknown function that takes into account the population activity, plus some external time-varying input.

The final equation, found at the end of the results, is similar to Equation 2, but written in the spirit of linear algebra. Given the observations in this paper, the authors conclude that just one matrix can capture the dynamics of neural activity underlying reaching, regardless of the properties of the reach.

How much you dig into this math depends a bit on the scope of your course—although dynamical systems is the theoretical backdrop, it's not central to understanding the concepts of this paper. If you or your students desire a deeper dive, the obvious choice is the book "Dynamical Systems in Neuroscience" (2010).

"To me, the real beauty of the original idea is that it made sense of a bunch of diverse facts, including otherwise-confusing features of single-neuron responses."

DIMENSIONALITY REDUCTION

Beyond interpreting the equations in the paper, there is also the big, flat, and rotating elephant in the room: dimensionality reduction. The paper, and this entire body of work, uses a version of principal component analysis (PCA) to project population activity into a lower-dimensional space where rotational features can be seen and measured. Unless you're teaching an upper-division computational neuroscience, modeling or mathematics course, you're probably not going to dive into the math behind PCA. That said, I think it's entirely possible to build an intuition for it using graphical explanations, such as the one below. Cunningham and Byron Yu, a neuroscientist at Carnegie Mellon University in Pittsburgh, Pennsylvania, also have a very accessible review article on the topic, and Neuromatch Academy has a dimensionality reduction tutorial with videos and code, as well as a tutorial on discrete dynamical systems.

THE SCIENTISTS BEHIND THE PAPER

Krishna Shenoy, the principal investigator behind this paper, made major contributions to the fields of neural dynamics and brain computer interfaces. He died of pancreatic cancer earlier this year. If you would like to learn more about him and his work more broadly, please read "A scientist's quest for better brain-computer interfaces opens a window on neural dynamics" or this obituary by Mark Churchland and Paul Nuyujukian.

Shenoy's team originally wrote the paper with the model first, but after leaving many people quite puzzled at their well-attended Society for Neuroscience conference poster, they reworked their story. Churchland and Cunningham rallied at the benches outside of the conference, staring over San Diego Bay, and decided they needed to reverse the order: unexpected results first, model second.

Now, 11 years after the paper's publication, Churchland reflects fondly on the ideas they presented and how they shifted the field:

"To me, the real beauty of the original idea is that it made sense of a bunch of diverse facts, including otherwise-confusing features of single-neuron responses. The hypothesis thus had appeal to me long before we were able to test that final prediction . . . At that time, it took a long time to get someone to understand both the essence and the appeal of the hypothesis. They were only willing to invest the time if they saw a novel result, presented first, that demanded an explanation."

This paper was, at its core, a team effort, and not only between the seven authors. Churchland also credits Larry Abbot and Evan Schaffer with helping to develop the ideas of the paper, as well as the approach to sharing it with other scientists.

FUTURE LESSONS

This paper spawned an entire field of scientists looking at the rotational dynamics in not just motor systems, but also cognitive ones. Ultimately, many still wonder why these dynamics are there, and what their role is. If you're looking for some food for thought on the purpose of these rotations, see my article "Discovery of rotational dynamics." This could also open up an interesting discussion about the links between structure, function and computation.

One great way for students to engage with this paper is to work with similar data themselves. This interactive notebook provides students with an introductory implementation of PCA using Python, and Yu has developed a problem set for MATLAB. Finally, Churchland and others have since built on this work, asking questions about what happens when trajectories of networks get tangled, and he spoke about it in 2020 in a lecture at the Bernstein Center for Computational Neuroscience in Munich, Germany.

From bench to bot: A scientist's guide to AI-powered writing

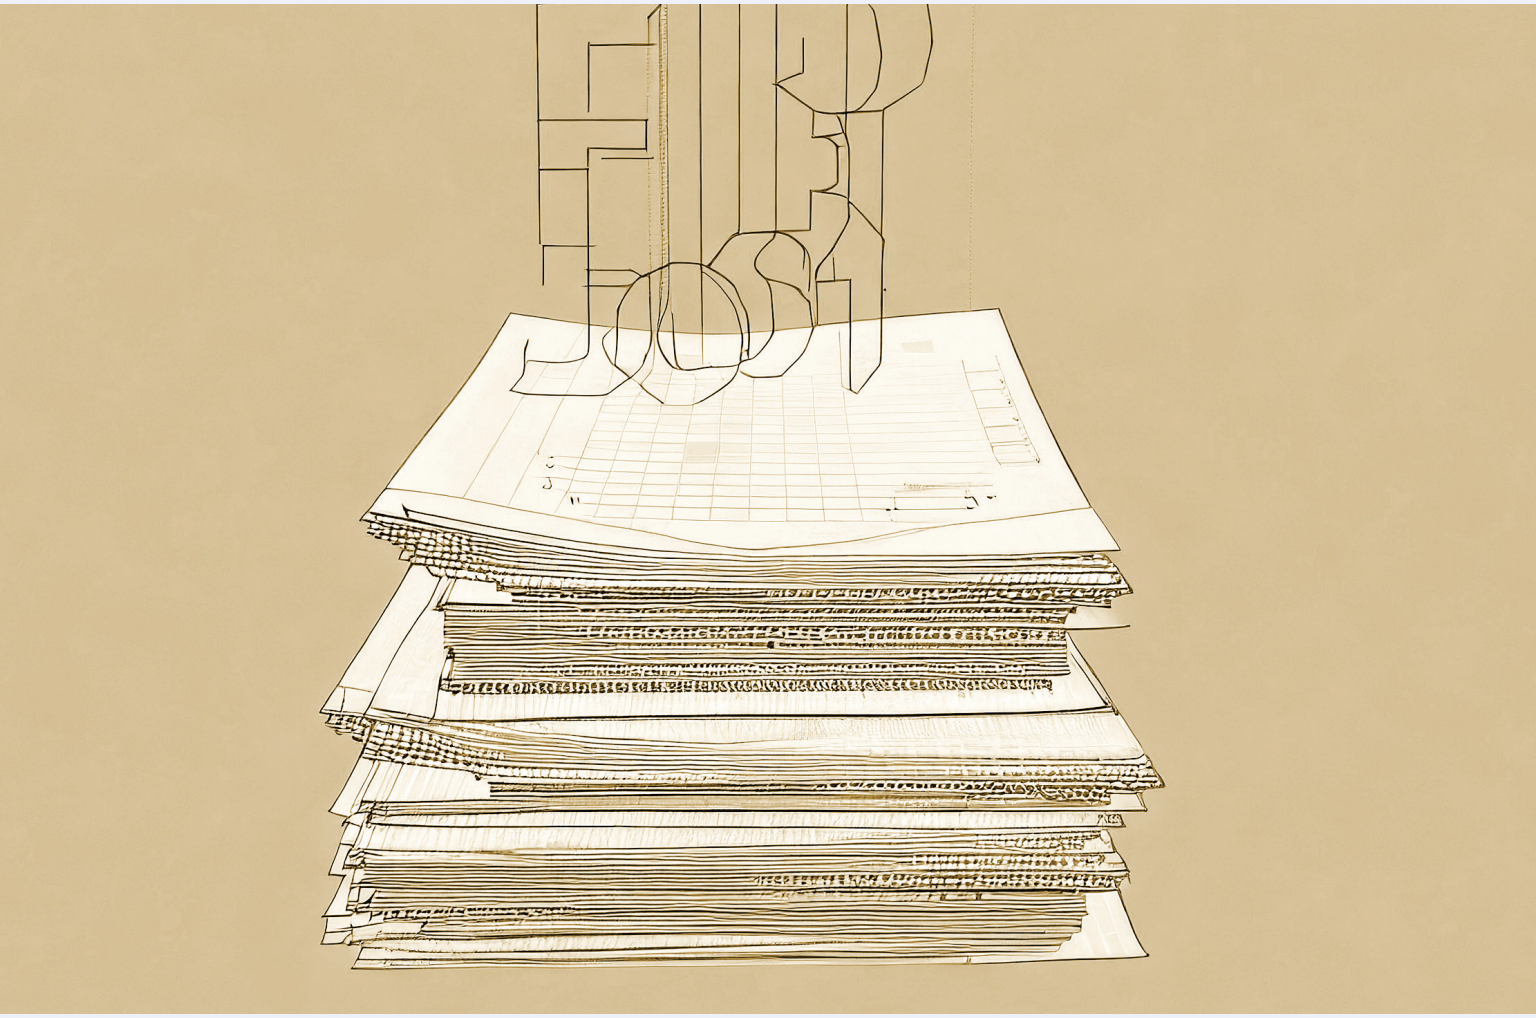

ILLUSTRATION BY REBECCA HORNE / ADOBE FIREFLY

I was initially skeptical of artificial-intelligence tools such as ChatGPT for scientific writing. But after months of using and teaching generative artificial intelligence, I have come to realize that it has a place in the scientific writer's tool kit, even if it can't write that grant for you from scratch.

As a scientist, you're a professional writer. You write grants to fund your research and papers to share your findings with the world. You work under deadline—and under pressure. You may not get paid per word or publish bestselling novels, but your livelihood depends on consistently producing quality writing on time.

That doesn't necessarily mean writing comes easily to you. So, when generative artificial-intelligence (AI) tools such as ChatGPT burst onto the scene late last year, perhaps you were allured by the in silico siren call. Could AI make this critical but challenging part of your job a little less stressful? Or perhaps you looked on with skepticism. Sure, these chatbots might be great for churning out marketing copy, but for scientists they're just a distraction, a computationally expensive way to produce semi-accurate, uninspired text.

As both a professional science writer and an instructor of scientific writing, I was curious about this new technology while also wary of its limitations and implications—the biases baked into this technology and the ethically questionable way it was built are causes for real concern. At the same time, I knew I couldn't pretend tools such as ChatGPT don't exist, because I'd need to be able to guide students, postdoc-

BY TIM REQUARTH, JOURNALIST AND LECTURER, NYU GROSSMAN SCHOOL OF MEDICINE

"Now is not the time for professional writers—scientists included—to bury their heads in the sand. In a decade, lacking proficiency in generative AI might be akin to not knowing how to use a search engine today."

toral researchers and principal investigators on using them—it's literally my job: For the past five years, I've been working full time at the Vilcek Institute of Graduate Biomedical Sciences at the NYU Grossman School of Medicine in New York City to develop and teach a scientific communication curriculum. If ChatGPT and other such innovations turned out to be powerful writing aids, I wouldn't want people with brilliant ideas but little GPT savvy to get left behind. If, on the other hand, AI-assisted writing was mostly hype, I'd need to show people who were already using it why it wasn't going to further their writing goals.

Since ChatGPT's launch in late 2022, I've immersed myself in the wonders and woes of generative AI. I've advised many scientists at my home institution and given talks and workshops on this fast-evolving field.

This monthly column will distill what I've learned—and am still learning—about how best to incorporate these tools into your writing process. I'll admit, I was initially skeptical of ChatGPT and similar tools for scientific writing. For intellectual work, writing struck me as too intimately linked to thinking to outsource it to a chatbot. "If people cannot write well, they cannot think well, and if they cannot think well, others will do their thinking for them," goes one of George Orwell's more famous quotes. I didn't trust ChatGPT to do the thinking for myself or those I teach.

But after months of using and teaching generative AI, I have come to realize that it has a place in the scientific writer's tool kit, even if it can't write that grant for you from scratch. And the technology is likely only going to get more powerful and become more ubiquitous. In my view, now is not the time for professional writers—scientists included—to bury their heads in the sand. In a decade, lacking proficiency in generative AI might be akin to not knowing how to use a search engine today.

In the columns that follow, I plan to offer a series of use cases that explore ways that this technology can make writing better, faster or easier. I also won't shy away from telling you when generative AI simply isn't up to the task. For each of these use cases, I'll draw on my years of experience teaching writing with traditional pedagogy and also offer practical AI-assisted workflows. We'll dive into the dark art of "prompt engineering," which is a fancy term for getting AI to do what you want it to. Equally important, we'll discuss strategies for "output curation," or how to ensure you always retain the authority to discern that AI responses are meeting standards of accuracy and rhetorical impact. In the end, I can't promise you that AI will solve all your writing troubles—and frankly I wouldn't trust anyone who claims it will—but I can promise you that you'll have more realistic expectations about what you can ask of AI, and what will still be asked of you.

USER BEWARE

Each column will appear with this warning, so heed it now: When exploring the use of AI, it's crucial to be aware that to incorporate it into our writing life is to navigate a minefield of possible dangers. AI can confidently produce convincing but inaccurate information (often called "hallucinations"), making it untrustworthy for factual queries, which means it is crucial that you have verification checkpoints in your workflow. Even accurate AI-generated content can be biased. It is well documented, for example, that social biases, such as racism and sexism, are embedded in and exacerbated by AI systems. AI may also recapitulate bias in subtler ways, such as by steering users toward established scientific ideas, which are more likely to be represented in the AI's training data.

Data-privacy concerns arise when using standard web interfaces, as user inputs can be adopted to train future AI models, though certain technical workarounds offer more protection. And at least one major journal (Science) and the U.S. National Institutes of Health have banned the use of AI for some purposes. Lastly, although generative AI generally does not pose a high risk of detectable plagiarism, that risk may increase for highly specialized content that is poorly represented in the training data (which might not be much of a concern for the typical user but could be a larger concern for the typical scientist). Some AI systems in development may overcome some of these problems, but none will be perfect. We'll discuss these and other issues at length as they arise.

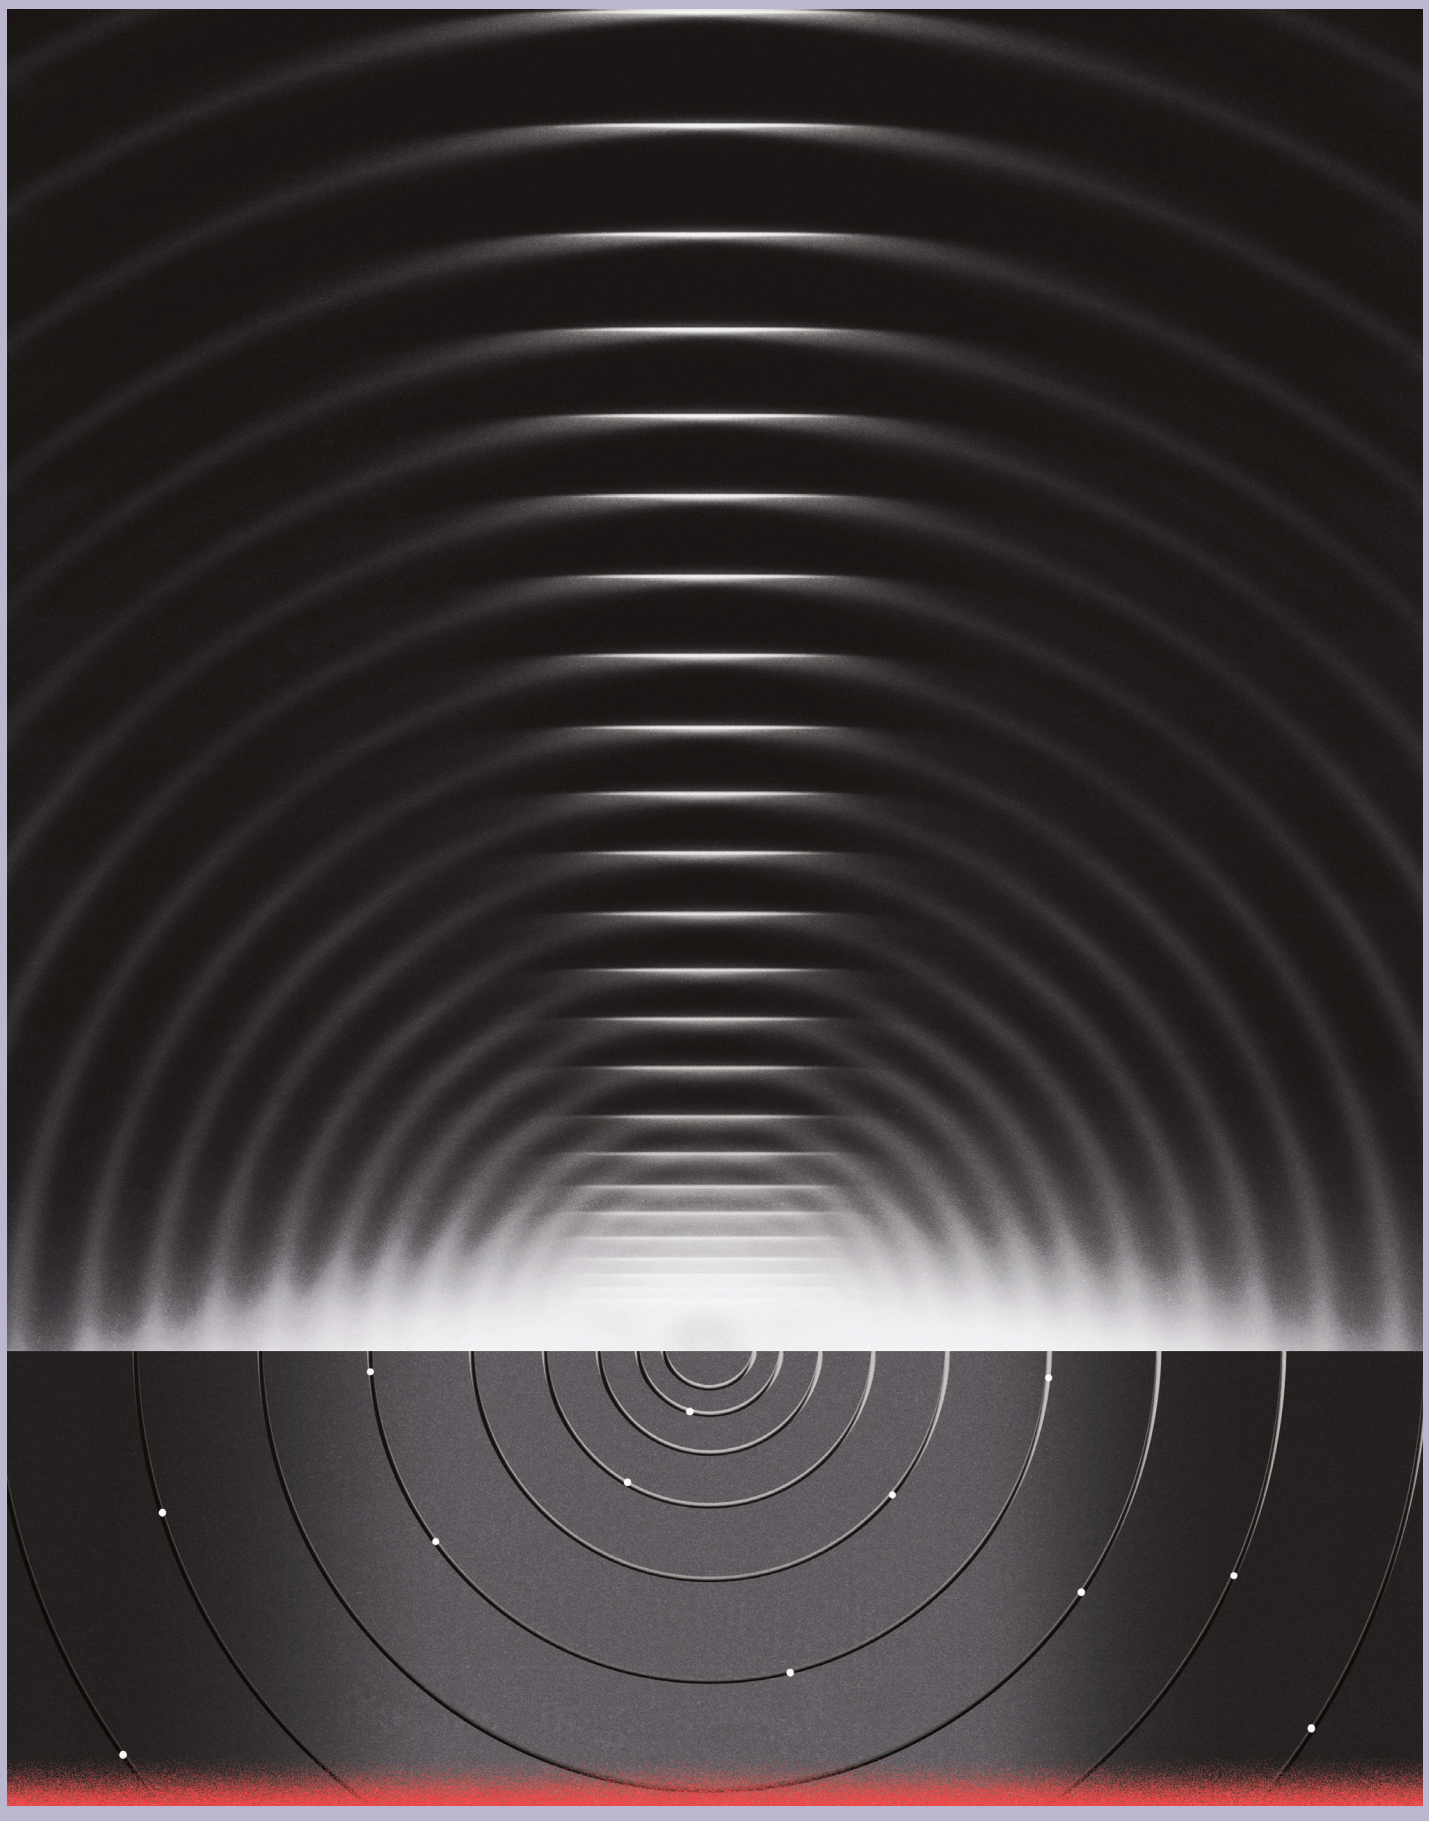

160 Fifth Avenue, 7th floor
New York, NY 10010

thetransmitter.org
@_TheTransmitter
@thetransmitter.bsky.social

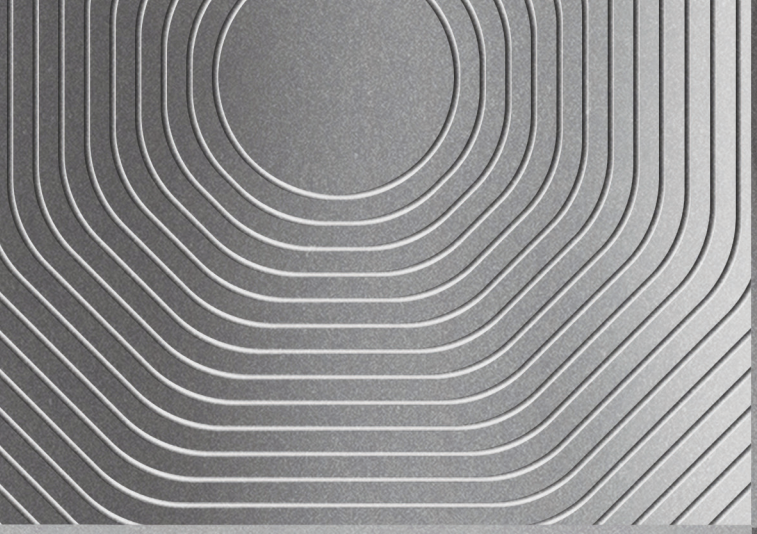
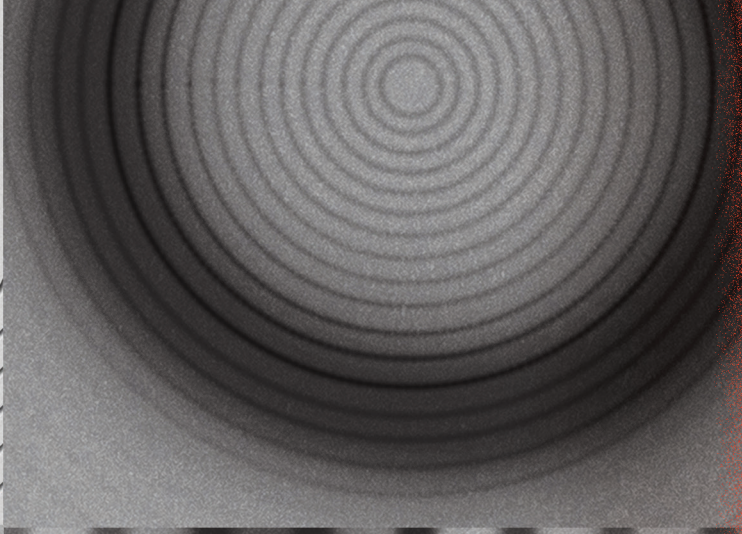
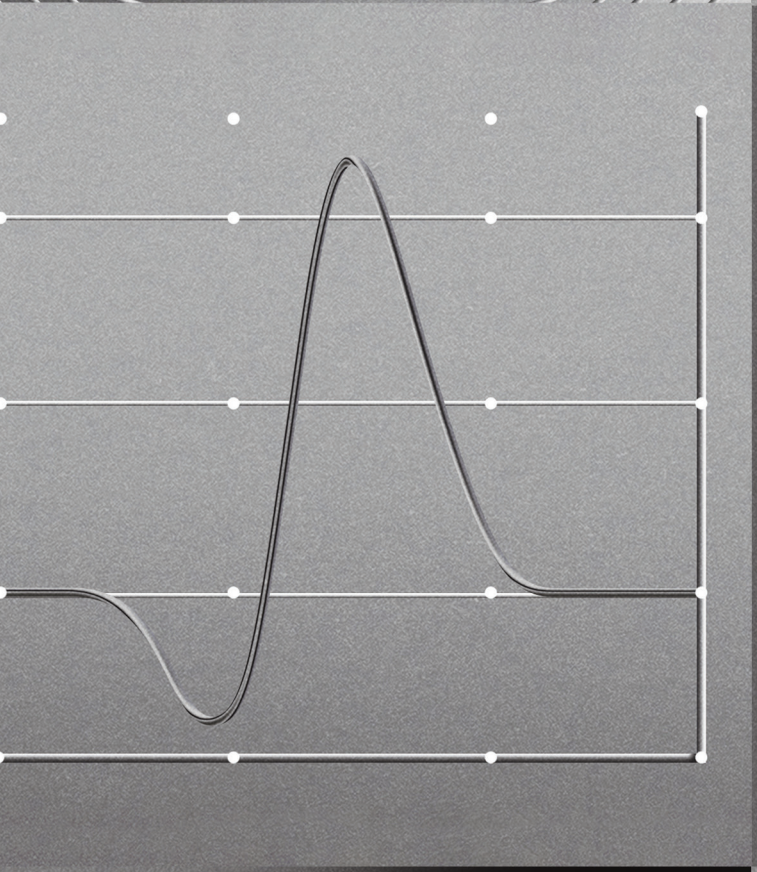
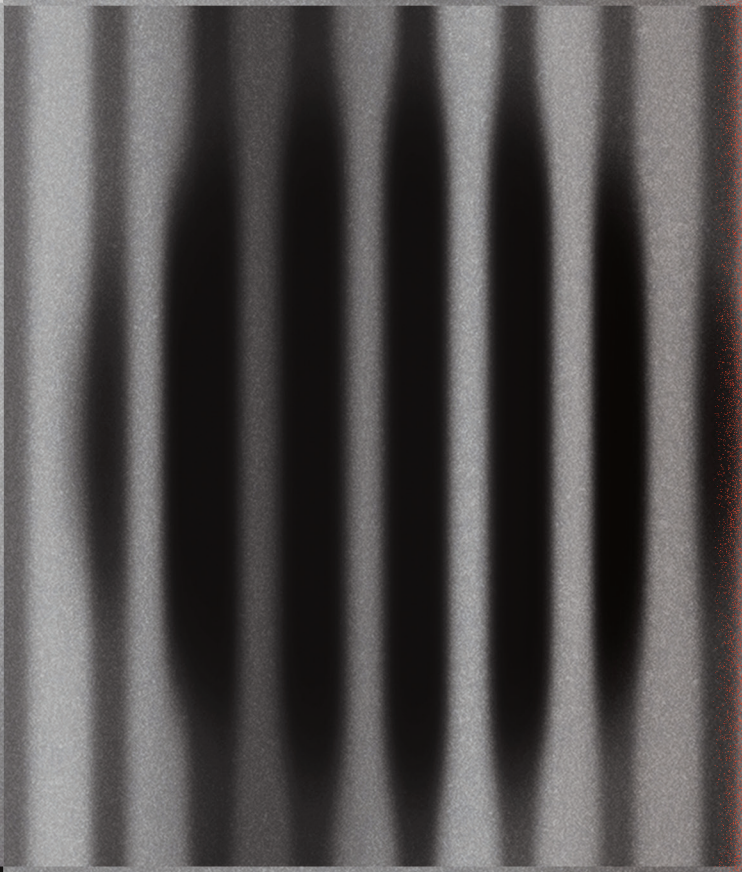

THE
TRANSMITTER